Source Unfoldings of Convex Polyhedra via Certain Closed Curves*

Jin-ichi Itoh† Joseph O'Rourke‡ Costin Vîlcu§

May 7, 2012

Abstract

We extend the notion of a source unfolding of a convex polyhedron $\mathcal{P}$ to be based on a closed polygonal curve Q in a particular class rather than based on a point. The class requires that Q "lives on a cone" to both sides; it includes simple, closed quasigeodesics. Cutting a particular subset of the cut locus of Q (in $\mathcal{P}$) leads to a non-overlapping unfolding of the polyhedron. This gives a new general method to unfold the surface of any convex polyhedron to a simple, planar polygon.

1 Introduction

Two general methods were known to unfold the surface $\mathcal{P}$ of any convex polyhedron to a simple polygon in the plane: the source unfolding and the star unfolding, both with respect to a point $x \in \mathcal{P}$; see, e.g., [DO07]. In [IOV10] we defined a third general method, the star unfolding with respect to any member of a class of simple, closed curves Q on $\mathcal{P}$, a class we called "quasigeodesic loops." Here we extend the source unfolding to be based on a different class of simple, closed, polygonal curves Q. The intersection of the two classes includes simple, closed quasigeodesics. Thus both the source and star unfolding are now generalized from points to these quasigeodesics, in all cases unfolding $\mathcal{P}$ to a simple, planar polygon.

Cut Locus. The *point source unfolding* cuts the *cut locus* $C_{\mathcal{P}}(x)$ of the point x: the closure of the set of all those points y to which there is more than one shortest path on $\mathcal{P}$ from x. The point source unfolding has been studied for polyhedral surfaces since [SS86] (where the cut locus is called the "ridge tree"). Our method also relies on the cut locus, but now the cut locus $C_{\mathcal{P}}(Q)$ with

*A preliminary abstract appeared in [IOV09].

†Dept. Math., Faculty Educ., Kumamoto Univ., Kumamoto 860-8555, Japan. j-itoh@kumamoto-u.ac.jp

‡Dept. Comput. Sci., Smith College, Northampton, MA 01063, USA. orourke@cs.smith.edu.

§Inst. Math. 'Simion Stoilow' Romanian Acad., P.O. Box 1-764, RO-014700 Bucharest, Romania. Costin.Vilcu@imar.ro.

respect to Q. The definition of $C_{\mathcal{P}}(Q)$ is the same: it is the closure of all points to which there is more than one shortest path from Q. Here it is analogous to the "medial axis" of a shape; indeed, the medial axis of a polygon is the cut locus of the polygon's boundary. As with the point source unfolding, our unfolding essentially cuts all of $C_{\mathcal{P}}(Q)$, but this statement needs to be qualified: we do not cut some segments of the cut locus (incident to Q), and we cut some additional segments not in the cut locus (again incident to Q).

Convex Curves and Quasigeodesics. Let p be a point on an oriented, simple, closed, polygonal curve Q on $\mathcal{P}$. Let $L(p)$ be the total surface angle incident to the left side of p, and $R(p)$ the angle to the right side. Q is a *convex curve* if $L(p) \leq \pi$ for all points of Q. Q is a *quasigeodesic* if $L(p) \leq \pi$ and $R(p) \leq \pi$ for all p, i.e., it is convex to both sides. Quasigeodesics, introduced by Alexandrov (e.g., [AZ67, p.16]) are the natural generalization of geodesics to polyhedral surfaces. A *quasigeodesic loop* has a single exceptional point at which the quasigeodesic angle condition to one side may not hold. We have the inclusions: {convex} $\supset$ {quasigeodesic loop} $\supset$ {quasigeodesic}.

The class of curves for which our source unfolding method works includes (a) convex curves that pass through at most one vertex of $\mathcal{P}$, and (b) quasigeodesics. The method does not (always) work for all quasigeodesic loops. Thus the class of curves for which the source and star unfolding methods work are not directly comparable, but they both include quasigeodesics.

Curves "Living on a Cone." The precise class of curves for which our source unfolding method works depends on the following notion. Let Q be (as before) an oriented, simple, closed, polygonal curve on $\mathcal{P}$. Let N be a vertex-free neighborhood of Q in $\mathcal{P}$ to the left of and bounded by Q. We say Q *lives on a cone* to its left if there exists a cone Λ and an N such that $Q \cup N$ may be embedded isometrically on Λ, enclosing the cone apex a. A *cone* is a developable surface with curvature zero everywhere except at one point, its apex a. We consider a planar polygon to be a cone with apex angle 2π, and a cylinder to be a cone with apex angle 0. The source unfolding described in this paper works for any curve Q that (a) lives on a cone to both sides (perhaps on different cones), and (b) such that each point of Q is "visible" from the apex a along a generator of the cone (a line through a lying in Λ). See Figure 1. We should remark that the cone on which a curve Q lives has no direct relationship (except in special cases) to the surface that results from extending the faces of $\mathcal{P}$ crossed by Q. The cones on which Q lives play a central role in our proof technique. Although a curve could live on many different cones, it is established in [OV11, Lem. 3] that the cone is uniquely determined to each side by $\mathcal{P}$. We also established [OV11, Thm. 3] that the classes of curves listed above live on cones in this sense. Indeed the set of curves enjoying the properties (a) and (b) above is wider than what we list here, but as the full class of curves that live on cones to both sides is not yet precisely delimited, we leave this issue aside.

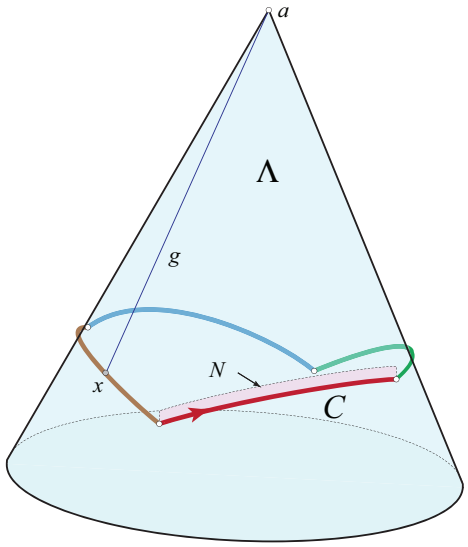

Figure 1: A 4-segment curve Q which lives on cone Λ to its left. A portion of N is shown, and a generator $g = ax$ is illustrated. (Adapted from [OV11].)

2 Preview of Algorithm

Assume we are given a curve Q satisfying our conditions: it is convex to one side, and it lives on cones to either side. We now describe the unfolding abstractly at a high level. First we need some notation.

Q divides $\mathcal{P}$ into two closed "halves," P_1 and P_2. We handle the two halves a bit differently. Let P_1 be the half to the convex side (say, the left side), and P_2 the half to the (possibly) reflex (i.e., nonconvex) right side. (If Q is a quasigeodesic, both halves are convex.) Let C_{P_i} be the portion of the cut locus $C_{\mathcal{P}}(Q)$ in each half P_i.

Cut all edges of C_{P_1} not incident to Q, and cut one additional precisely selected segment from C_{P_1} to Q. To the (possibly) nonconvex side P_2, cut all of C_{P_2}, including those edges incident to Q. In addition, we cut shortest path segments from C_{P_2} to any remaining reflex vertices of Q. See Figure 2. The result, we will show, is an unfolding of $\mathcal{P}$ to a simple, planar polygon.

We do not concentrate in this paper on algorithmic complexity issues, which will only be touched upon in Section 7. But the described procedure is a definite, finite algorithm which works for any Q in the appropriate class.

The unfolding procedure is best viewed as unfolding each half separately, and then gluing them together along Q, as the examples below will emphasize.

Examples. Throughout this paper we will use three example polyhedra $\mathcal{P}$ to illustrate concepts. Here we introduce two; the third will be described in Section 3.

Example 1. The first example is a regular tetrahedron with a convex curve

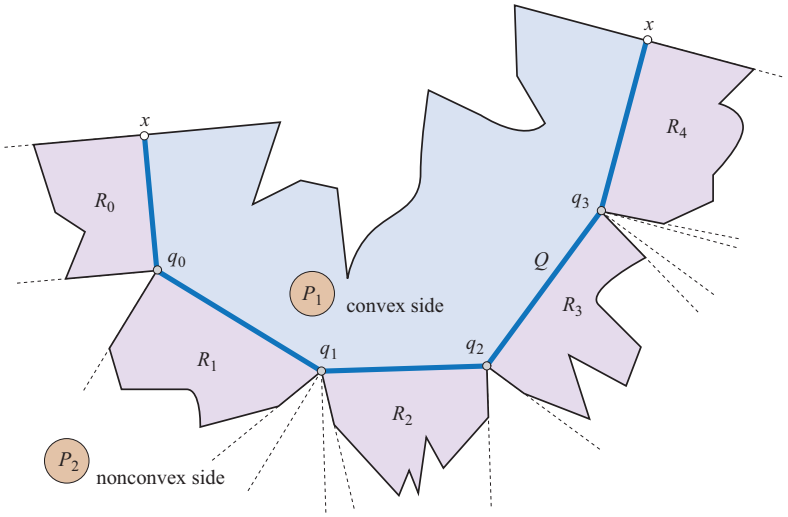

Figure 2: An abstract depiction of the full source unfolding for a convex curve Q. Here q_1 and q_3 are reflex vertices of Q to the P_2 side. (The figure is not metrically accurate.)

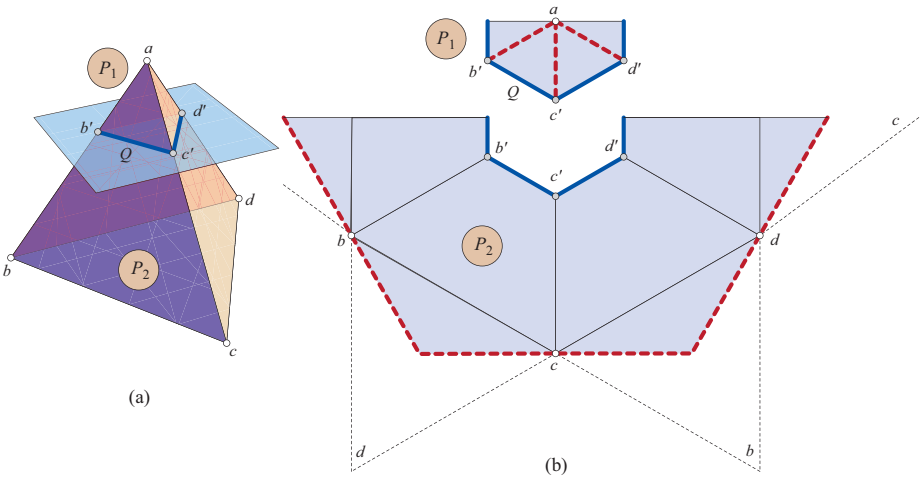

Figure 3: (a) Regular tetrahedron sliced at $\frac{2}{3}$'s of its altitude. (b) Unfoldings of P_i, with cut loci C_{P_i} drawn dashed. C_{P_2} partitions the bottom face $\triangle bcd$ into three congruent triangles meeting at its centroid.

parallel to the base; see Figure 3(a), where $Q = (b', c', d')$. Q lives on the same cone to each side, the cone determined by the lateral faces of the tetrahedron. Note that the angle at each corner of Q is convex ($\frac{2}{3}\pi$) to one side and reflex ($\frac{4}{3}\pi$) to the other side. One of our main results, Theorem 1, shows that each half unfolds separately without overlap, as illustrated in Figure 3(b). A second main result, Theorem 2, shows that the two halves may be joined to one non-overlapping piece, in this case producing a trapezoid. This is an atypical case in many respects, but will be useful for that reason to illustrate degeneracies. For example, notice that no part of C_{P_1} is cut in this example, because all segments of that cut locus are incident to Q.

Example 2. Our second example is more generic: a cube twice truncated, with Q the particular quasigeodesic shown in Figure 4(a). The angles at the vertices of $Q = (v_0, v_1, v_7, v_{10})$ within P_1 are, respectively, ($\frac{3}{4}\pi$, $\frac{1}{2}\pi$, $\frac{3}{4}\pi$, $\frac{1}{2}\pi$), and the angles at those vertices within P_2 are ($\frac{3}{4}\pi, \pi, \frac{3}{4}\pi, \pi$). Because all of these angles are at most π, Q is convex to both sides and so a quasigeodesic. The cone on which Q lives to the P_2-side is evident: its apex is the cube corner truncated. The cone on which Q lives to the P_1-side is not evident; it will be described later (in Figure 6). For Q a simple closed quasigeodesic, the cut loci C_{P_i} are each a single tree with each edge a (geodesic) segment, as illustrated in Figure 4(b,c). Theorem 2 leads to the unfolding shown in Figure 4(d).

Our third example will illustrate that the edges of the cut locus can also be parabolic arcs.

3 Cut Locus of Q

The proof that the described source unfolding avoids overlap relies on two key ingredients: the structure of the cut loci in each half of $\mathcal{P}$, and the cones on which Q lives. Here we focus on the cut loci.

A *vertex of* Q is a point q with an angle at q in P_i different from π; note this definition depends on the half P_i. Thus if Q is a geodesic, it has no vertices at all to either side. It will be useful to distinguish two varieties of these vertices: a *convex vertex* has angle less than π in P_i, and a *reflex vertex* has angle greater than π. Note the meaning of these is interchanged when looking from P_1 compared to looking from P_2 (although it is possible that an angle differs from π from one side and is equal to π from the other, e.g., v_{10} in Figure 4). Let $q_0, q_1, \ldots, q_k$ be the vertices of Q in some circular order, with respect to the half P_i. (Although these vertices depend on P_i, we opt not to subscript with i to ease notation.)

The cut locus C_{P_i} is a tree whose leaves span the vertices of P_i, including the convex vertices of Q, which must be leaves of C_{P_i}. (That C_{P_i} is a tree is well-known in Riemannian geometry, e.g., see [Thu98, p. 539]. This can also be seen from the fact that each half P_i has a finite intrinsic diameter, and so shortest paths from Q are finite in length and each ends at a cut point.) Each non-leaf point p of C_{P_i} has at least two shortest paths to Q; leaves have precisely one shortest path to Q, or might lie on Q. Each shortest path from $p \in C_{P_i}$

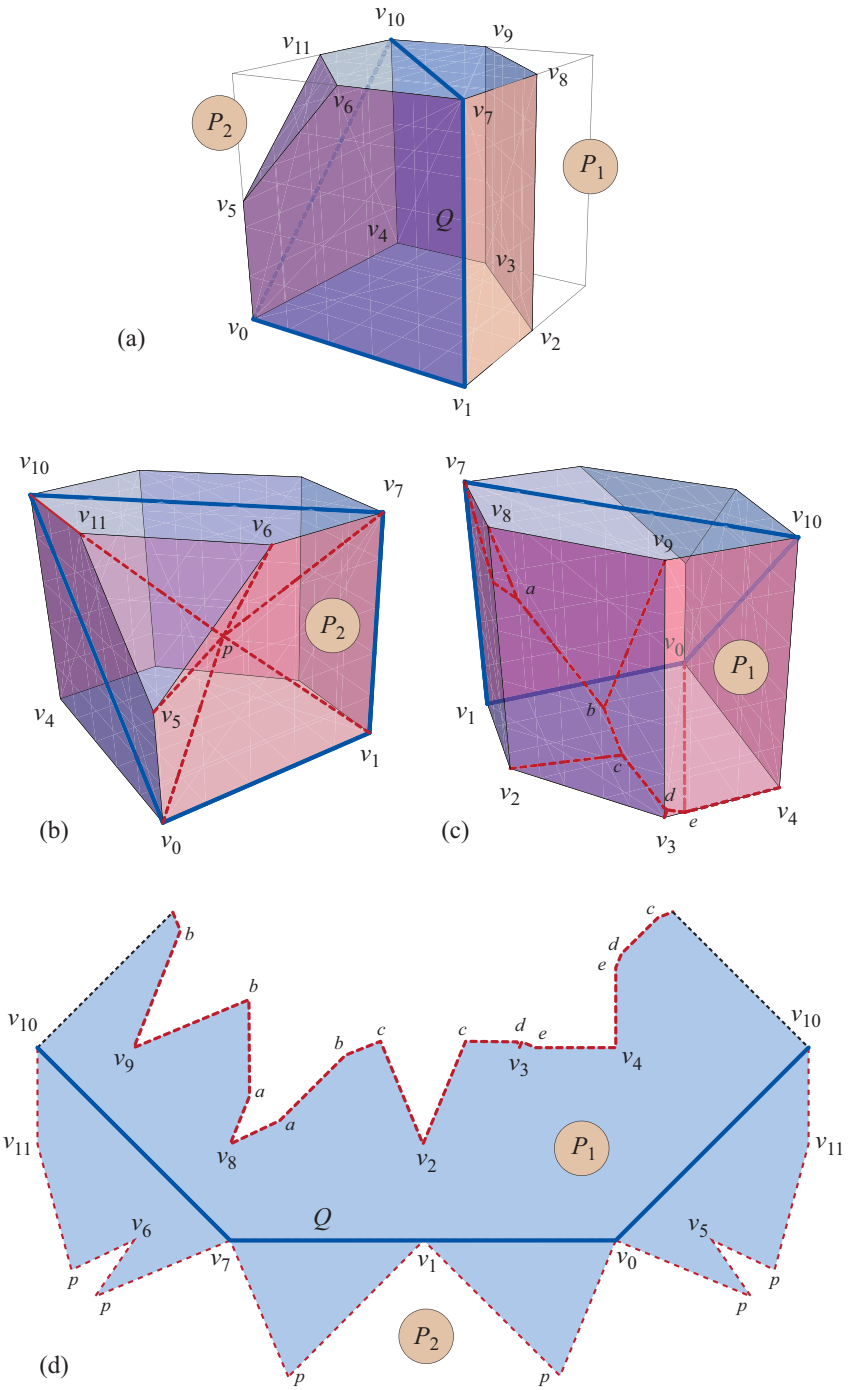

Figure 4: (a) Truncated cube and quasigeodesic $Q = (v_0, v_1, v_7, v_{10})$. (b,c) Views of P_i and cut loci C_{P_i} (dashed). (d) Unfolding to a simple polygon.

to Q is called a *projection* to Q. The segment of C_{P_i} incident to each convex vertex of Q bisects the angle there.

We now review our three examples to illustrate these relationships. In Figure 3, C_{P_1} is a tree spanning the apex a and the three vertices $\{b', c', d'\}$ of Q. C_{P_2} is a tree in the bottom face spanning its three vertices $\{b, c, d\}$ (but not touching Q).

Figures 4(b,c) show that C_{P_i} are each trees spanning the vertices of P_i and touching Q: at its four vertices in P_2, and its two vertices in P_1 (the angles at v_1 and at v_{10} are π in P_1).

Example 3. Figure 5 shows an example that illustrates several aspects not present in the other two examples. First, Q is neither a convex curve nor a quasigeodesic, but it nevertheless lives on a cone to both sides, namely, the cone determined by the pyramid's lateral sides. Thus our method applies to this Q. Second, C_{P_1} is a tree spanning the apex a and the two convex vertices of Q to that side, $\{b', d'\}$. C_{P_1} includes four parabolic arcs: one generated by the edge $b'e'$ and the reflex vertex c'_1, one generated by the edge c'_1c' and the reflex vertex e', and two more symmetrically placed arcs. C_{P_2} is a tree that includes an `X` on the bottom face spanning $\{b, c, d, e\}$. In general, parabolic arcs arise as arcs at equal distance to a reflex vertex and an edge.

4 Cut Loci on Cones

The essence of our proof, that the unfolding of each half P_i avoids overlap (Theorem 1), is that pieces of P_i embed into the cone Λ_i, and that Λ_i develops without overlap in the plane when cut along a generator. A key to our approach is to define these "pieces" of P_i (in Sec. 5) by comparing the cut locus C_{P_i} on P_i with the cut locus C_{Λ_i} on the cone Λ_i on which Q lives to the P_i-side. The cone Λ_i and surface P_i share the same boundary Q, and by construction, the same angles occur along Q. Therefore, Q has the same vertices in both Λ_i and P_i. Thus C_{P_i} and C_{Λ_i} both have the same set of leaves touching Q, but of course in general they differ in the interior of the surfaces.

To one of the two sides, the cone may be unbounded (established in [OV11]). This occurs in both Figure 3 and Figure 5. For the tetrahedron (Figure 3), Λ_2 is unbounded, and C_{Λ_2} is empty. For the pyramid (Figure 5), Λ_2 is again unbounded, and C_{Λ_2} is in this case a forest, including two halfline branches to infinity, one from e' through e, and another from the `Y`-junction below c' through c. The cut locus can only be a forest (as opposed to a tree) to an unbounded side. This is because an unbounded surface Λ_i contains infinite-length shortest paths without cut points, which then separate C_{Λ_i} into components, as in the pyramid example. Only one side may be unbounded, unless the cone is a cylinder.

In the truncated cube example, both cones are bounded (because the quasigeodesic Q is convex to both sides), and both C_{Λ_1} and C_{Λ_2} are trees; see Figure 6. Note that the cut locus extends to the apex a_i of Λ_i in both instances, which it must because the cut locus includes all vertices.

C_{Λ_i} is determined by any small neighborhood of Q to the P_i-side, because

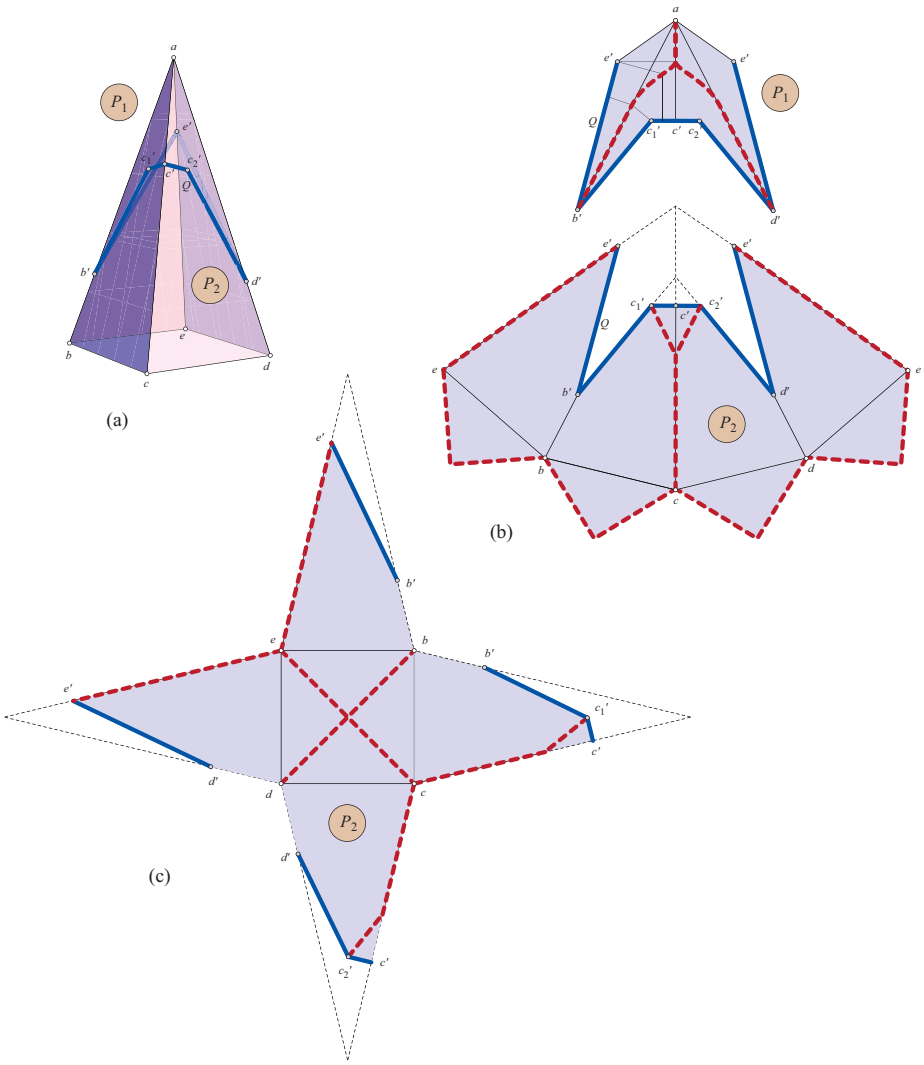

Figure 5: (a) A square-based right pyramid, with altitude twice the base side length. $|bb'| = |dd'| = |ae'| = \frac{1}{4}|ab|$, and $|b'c_1'| = \frac{3}{4}|b'e'|$. (b) Source unfolding of the two halves P_1 and P_2. (c) A view of P_2 from the bottom of the pyramid.

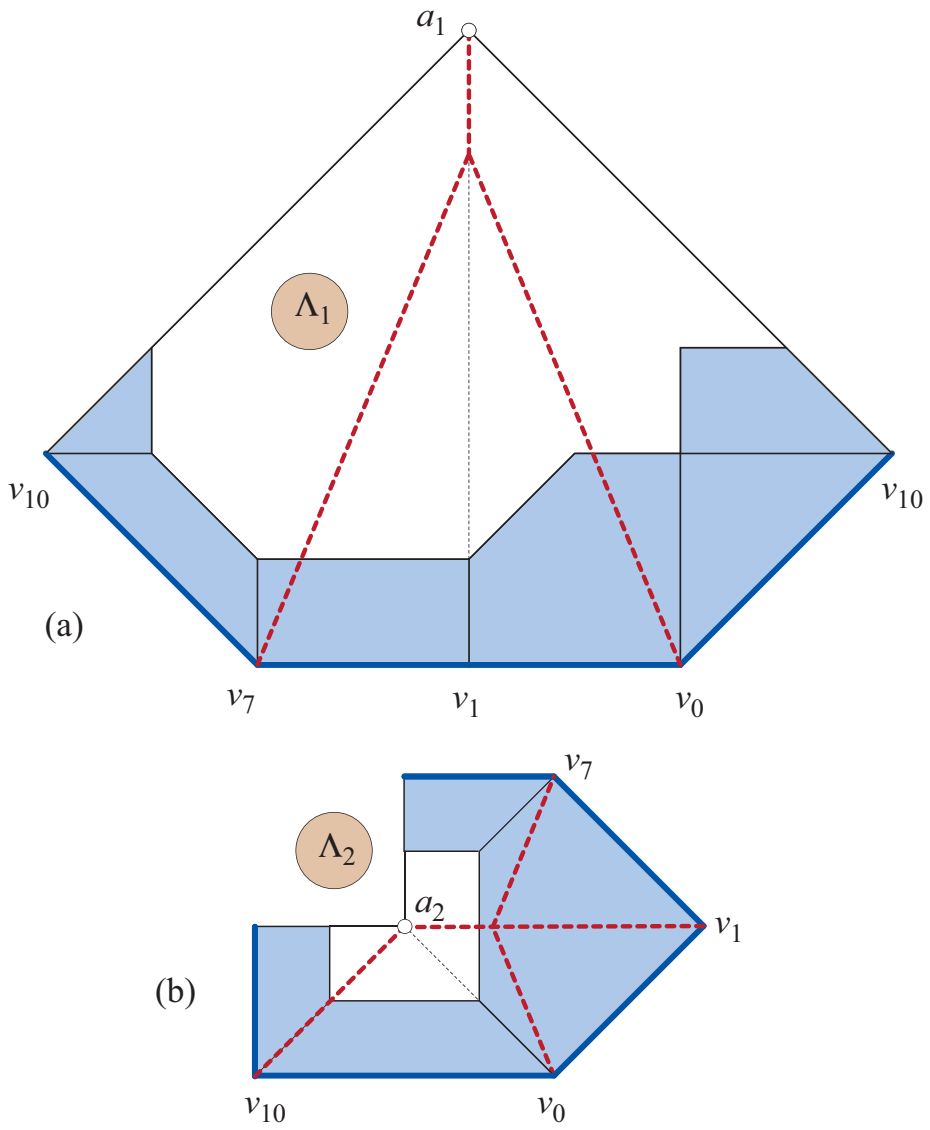

Figure 6: Developed cones Λ_1 and Λ_2 for $\mathcal{P}$ in Fig. 4(b,c), with C_{Λ_1} and C_{Λ_2} indicated.

Λ_i is so determined (see again [OV11]). And because every leaf of C_{Λ_i} is also a leaf of C_{P_i}, the bisecting property of cut loci (each edge of the cut locus incident to a vertex $v \in Q$ bisects the angle of P_i at v) implies that small neighborhoods of the leaves of C_{Λ_i} are included in C_{P_i}. In other words, the edges of C_{P_i} and C_{Λ_i} issuing from vertices of Q coincide until they hit a vertex of C_{P_i} or a vertex of P_i. We use this property in the proof of Lemma 1 below.

We now turn to defining the "pieces" of P_i that embed in Λ_i.

5 Peels & Subpeels: Embedding in the Cone

Let $u_0, u_1, \ldots, u_m$ be the vertices (leaves and junction points) of C_{P_i}, following a circular ordering of all of their their projections to Q. Note that this ordering is unambiguous even though some points have multiple equal-length projections to Q, because these projections never cross. Therefore, those leaves and junction points of C_{P_i} appear several times in the sequence $u_0, u_1, \ldots, u_m$, each as many times as its number of projections. Let (u_j, u_k) be two consecutive leaves of C_{P_i}, and (u'_j, u'_k) corresponding consecutive projections onto Q. The *peel* $\alpha_{P_i}(u_j, u_k)$ is the closed flat region of P_i bounded by the two projection paths $u_j u'_j$, $u_k u'_k$, the subpath Q_{ij} of Q from u'_j to u'_k, and the unique path in C_{P_i} connecting u_j to u_k, such that $\alpha_{P_i}(u_j, u_k)$ contains no leaf of C_{P_i}. Each peel $\alpha_{P_i}(u_j, u_k)$ is isometric to a planar convex polygon if Q_{ij} is convex to the left; otherwise, $\alpha_{P_i}(u_j, u_k)$ can be decomposed into the union of planar convex polygons and triangles whose base is a parabolic arc and whose vertex opposite to that side is a reflex vertex of Q.

Between two consecutive leaves u_j and u_k are the vertices along the tree path, $u_j, u_{j+1}, \ldots, u_{k-1}, u_k$. Each of these delimits a *subpeel* of $\alpha_{P_i}(u_j, u_k)$, partitioning the peel along the projection segments $u_l u'_l$, $j < l < k$.

The notion of peel and subpeel can be defined analogously for Λ_i. If C_{Λ_i} has at most one component, everything defined above for C_{P_i} holds exactly as stated. Assume now that C_{Λ_i} has at least two components. In this case, all leaves of C_{Λ_i} belong to Q, and we consider them as projection points of themselves. Again we consider all projections of the junction points of C_{Λ_i}. We take the circular order along Q of all these projection points. If two consecutive projections correspond to points on the same component of C_{Λ_i}, we use the definition of peels and subpeels given above.

Now assume two consecutive projections u_j and u_k correspond to points on different components of C_{Λ_i}. Then the peel $\alpha_{\Lambda_i}(u_j, u_k)$ is the closed flat region of Λ_i bounded by the two projection paths $u_j u'_j$, $u_k u'_k$, the subpath Q_{ij} of Q from u'_j to u'_k , and the two paths in C_{Λ_i} from u_j, respectively u_k, to infinity, such that $\alpha_{P_i}(u_j, u_k)$ contains no leaf of C_{P_i}. Here we note that each component of C_{Λ_i} has precisely one arc going to infinity, so the definition above is correct. The definition for subpeels is analogous: either we get bounded polygons, or unbounded polygons determined by arcs to infinity in different components of C_{Λ_i}.

These peels can be decomposed into the union of (a) planar rectangular

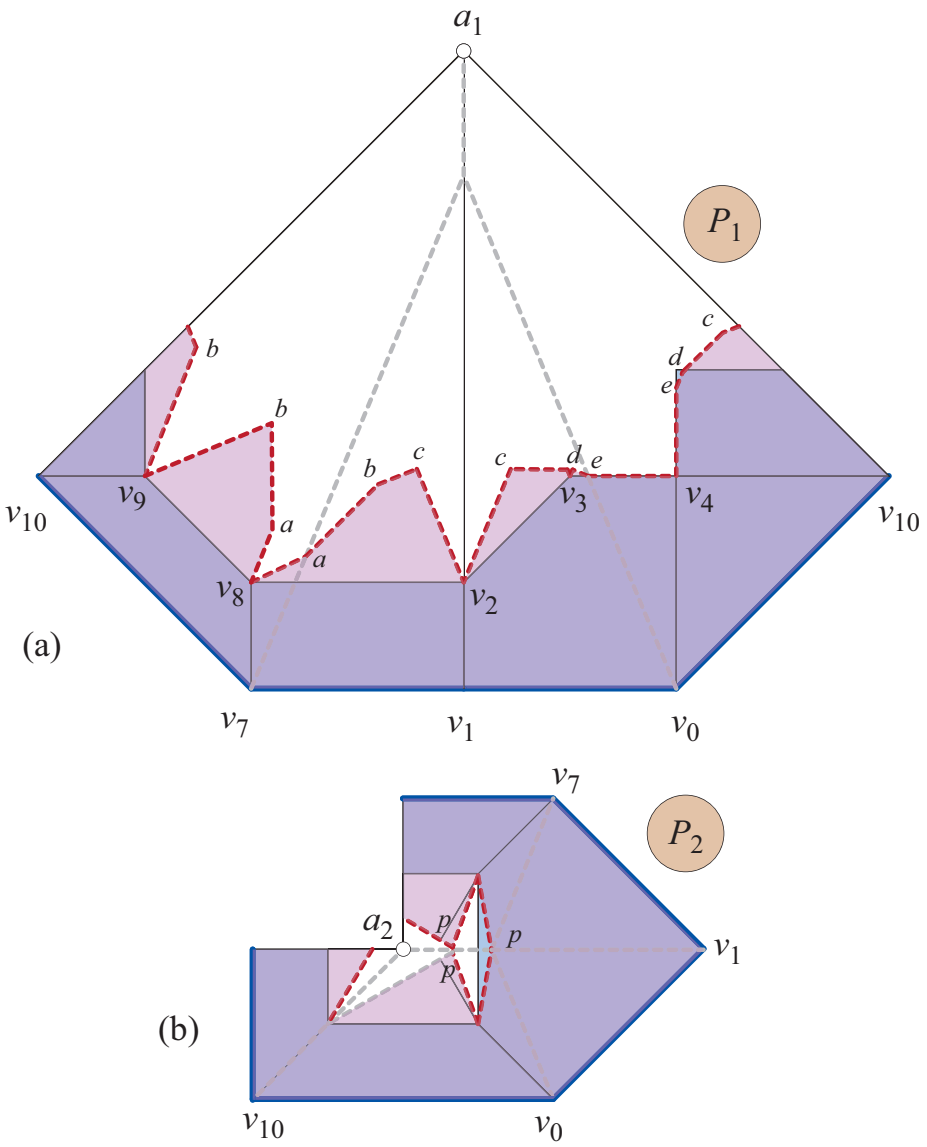

Figure 7: (a) Nesting of cut P_1 within the unfolded Λ_1 of Fig. 6(a). (b) Similar nesting of P_2 within Λ_2 of Fig. 6(b).

trapezoids, with one side either a line-segment, or a parabolic arc, or "an arc at infinity"; and (b) triangles whose base is a parabolic arc (possibly at infinity) and whose vertex opposite to that side is a reflex vertex of Q.

We now prove the central technical result, Lemma 1, which establishes the embedding of P_i into Λ_i. This lemma shows that the peels of P_i nest inside the peels of Λ_i. At one spot in the argument, we need a specialized lemma, which we invoke (Lemma 2) before proving it. As a consequence of the nesting lemma, embedding of the half-surfaces into their cones follows, as summarized in Lemma 3.

Lemma 1 (Peel Nesting) *Each subpeel β_{P_i} of P_i is isometric to a region of a subpeel β_{Λ_i} of Λ_i. The union of the subpeels in one peel α_{P_i} of P_i is non-overlapping in some α_{Λ_i}, and thus each peel α_{P_i} is nested inside a peel α_{Λ_i} of Λ_i.*

Proof: Let Γ_{P_i} be the directed curve that traces around the maximal subtree of C_{P_i} disjoint from Q, i.e., around the cut locus minus the edges incident to Q. Γ_{P_i} is an Eulerian tour of this subtree, tracing its edges twice, once from each side. Thus each non-leaf point x of $C_{P_i} \setminus Q$ has at least two images in Γ_{P_i}, and a junction point x of C_{P_i} has $\deg(x) \geq 3$ images in Γ_{P_i}. We define Γ_{Λ_i} to be the "image" of Γ_{P_i} on Λ_i: an isometric tracing with the same angles, on Λ_i. We are going to track a variable point x_t on Γ_{Λ_i} inside Λ_i, and analyze how x_t interacts with C_{Λ_i}. The crux of the proof analyzes what happens when x_t might leave its peel α_{Λ_i}.

We will illustrate the proof with the portion of Γ_{Λ_1} for P_1 of our truncated cube example shown in Figure 8.

To initiate the analysis, let $x_0 \in \Gamma_{\Lambda_i}$ be any point in the interior of a peel α_{Λ_i}. We need to argue for the existence of such a point. We will choose a leaf of C_{P_i}, but it must be chosen with some care. First, if $C_{P_i} = C_{\Lambda_i}$, then the claim of the lemma is trivial. So assume the cut loci differ. Delete from C_{P_i} all branches in common with C_{Λ_i}. By "branches" here we mean subtrees of C_{P_i} whose removal does not disconnect C_{P_i}; so what remains is still a tree.

Let $v \in P$ be a leaf of this reduced tree. As we argued above, the image in Λ_i of a small neighborhood of v in P_i remains included in Λ_i. So we may take $x_0 = v$. Vertex v_8 in Figure 8 could serve as x_0.

Now we move a point x_t along Γ_{Λ_i} continuously from x_0. As long as x_t remains inside the peel α_{Λ_i}, the subpeel β_{P_i} to which x_t belongs remains included in peel α_{Λ_i}. Adjacent subpeels β_{P_i} and β'_{P_i} share a side orthogonal to Q in α_{Λ_i}, and so they do not overlap one another.

Now assume that x_t reaches a point $x_t = y \in C_{\Lambda_i}$. If Γ_{Λ_i} touches but does not cross C_{Λ_i}, then the subpeel β_{P_i} to which x_t belongs, or the adjacent subpeel β'_{P_i} into which x_t is moving, remains included in α_{Λ_i}, and there is nothing to prove. We should note that, in general, we cannot conclude that touching-but-not-crossing C_{Λ_i} necessarily implies that y is a junction of C_{P_i}. It could be that y is a leaf of C_{Λ_i}, which is another instance of touching-but-not-crossing, and again nesting remains clearly true. So henceforth we assume y is interior to C_{Λ_i}.

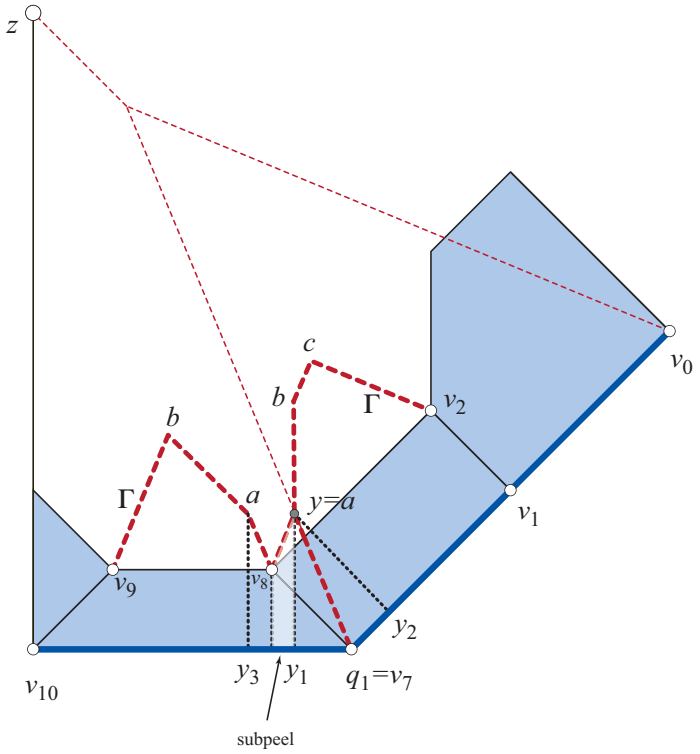

Figure 8: Γ_{Λ_1} is the dashed curve, directed left-to-right (from v_9 to v_8). The current subpeel is that bounded on the left by v_8. x_t crosses C_{Λ_1} at y, which coincides with the junction point a of C_{P_1}; see Fig. 4(c). The three segments are the projections of $y{=}a$ on P_1; y_1 and y_2 are the only two projections of y on Λ_1. Γ_{Λ_1} enters the new subpeel $\beta_{\Lambda_1}(v_7, y_2)$ and the new peel is $\alpha_{\Lambda_1}(v_7, v_0)$.

Now we consider the situation when x_t crosses C_{Λ_i} at y, as it does at $y{=}a$ in Figure 8. This is exactly when the claim of the theorem might be false, for the subpeel β_{P_i} bounded by Γ_{Λ_i} then extends beyond the peel α_{Λ_i}. But Lemma 2 below establishes that y must be (the image of) a junction point of C_{P_i} (and note that a is a junction in our example). This means that the subpeel β_{P_i} ends at y, and Γ_{Λ_i} enters a new subpeel β'_{P_i} of a new peel α'_{Λ_i} of Λ_i. Thus the subpeel nesting claim of the lemma is established. The peel nesting property will be obtained following the proof of Lemma 2. □ □

We complete the proof of Lemma 1 with a technical lemma that was invoked above. We use $\delta_S(x, y)$ to represent the distance function on surface S: the length of a shortest path on S between x and y. We will employ the Alexandrov-Toponogov Comparison Theorem in two versions, Lemmas 5 and 6 in the Appendix.

Lemma 2 (Junction Crossing) *When Γ_{Λ_i} crosses C_{Λ_i} at $x_t = y$, y is (the image of) a junction point of C_{P_i}.*

Proof: Because Γ_{Λ_i} crosses C_{Λ_i} at y, y has at least two projections to Q on Λ_i (because every interior point of C_{Λ_i} has at least two projections). One of the two segments, say yy_1, also is a segment on P_i, because up to this point $x_t{=}y$ in our tracing, the peel α_{P_i} containing x_t is nested in α_{Λ_i}. So the length of these segments is the same on P_i and on Λ_i: $\delta_{P_i}(y, y_1) = \delta_{\Lambda_i}(y, y_1)$.

Let y_2 be the "next" projection of y to Q on Λ_i, i.e., there is no other Λ_i projection of y between y_1 and y_2 along Q. See Figure 8, which illustrates these two Λ_i projections. We aim to establish that $\delta_{P_i}(y, y_2) \leq \delta_{\Lambda_i}(y, y_2)$, which will imply equality (because all projection segments from y have the same length on P_i, and the same length on Λ_i—details below). This will imply that the regions on P_i and on Λ_i determined by $\{y, y_1, y_2\}$ are isometric. Thus y would have two projection segments yy_1 and yy_2 on P_i. But these two segments derive from $x_t \in \Gamma_{\Lambda_i}$ corresponding to a point $y \in C_{P_i}$, and there is a second point of Γ_{Λ_i} that corresponds to the same y, on "the other side" of the cut locus. Thus there must be a third projection from y to y_3 on P_i, which implies that y is a junction of C_{P_i}. This situation is illustrated in Figure 8, where two points along Γ_{Λ_i} derive from "different sides" of a.

Now we prove the claim that $\delta_{P_i}(y, y_2) \leq \delta_{\Lambda_i}(y, y_2)$. Consider the vertices $q_1, \ldots, q_k$ between y_1 and y_2, ordered from y_1 to y_2. (In Figure 8 there is just one such vertex q_1; see Figure 9 for a generic example.) Then the right triangle $\triangle yy_1q_1$ is flat on Λ_i, and non-negatively curved on P_i. Hence by the Alexandrov-Toponogov Comparison Theorem (Lemma 6), the hypotenuse is no longer on P_i, $\delta_{P_i}(y, q_1) \leq \delta_{\Lambda_i}(y, q_1)$, and the angle at q_1 is at least as large on P_i (Lemma 5), $\angle_{P_i}(y, q_1, y_1) \geq \angle_{\Lambda_i}(y, q_1, y_1)$. This angle inequality implies that $\angle_{P_i}(y, q_1, q_2) \leq \angle_{\Lambda_i}(y, q_1, q_2)$; see Figure 9. Continuing with the same logic, $\delta_{P_i}(y, q_2) \leq \delta_{\Lambda_i}(y, q_2)$, and $\angle_{P_i}(y, q_2, q_1) \geq \angle_{\Lambda_i}(y, q_2, q_1)$, so $\angle_{P_i}(y, q_2, q_3) \leq \angle_{\Lambda_i}(y, q_2, q_3)$. This leads (by induction) to the conclusion that the last distance is no longer on P_i: $\delta_{P_i}(y, y_2) \leq \delta_{\Lambda_i}(y, y_2)$, which is exactly what we aimed to establish. We now show that there must be equality here.

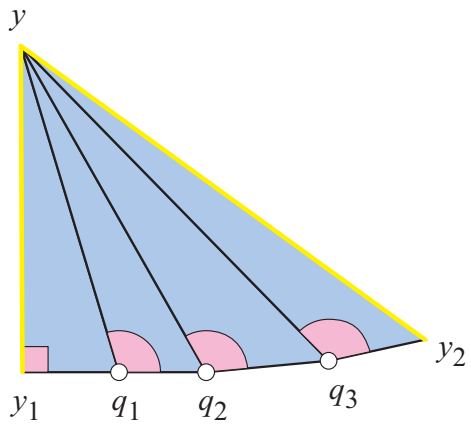

Figure 9: Applying the Alexandrov-Toponogov Comparison Theorem to the sequence of triangles apexed at y along Q from y_1 to y_2.

Suppose instead some inequality in the chain of reasoning above were strict. This leads to $\delta_{P_i}(y, y_2) < \delta_{\Lambda_i}(y, y_2)$. But we already know that $\delta_{\Lambda_i}(y, y_2) = \delta_{\Lambda_i}(y, y_1) = \delta_{P_i}(y, y_1)$, and so $\delta_{P_i}(y, y_2) < \delta_{P_i}(y, y_1)$. But this is a contradiction, because all projections from y to Q on P_i have the same minimal length.

Thus our conclusion above that there must be a third projection from y to y_3 on P_i follows, which means that that y is a junction of C_{P_i}. □ □

Having established the subpeel nesting property, we obtain next the same for peels. Outside the isometric regions of P_i and Λ_i that correspond to the external common subtrees, Γ_{Λ_i} is in one-to-one correspondence with what remains from Q (because of the projection onto along segments). I.e., each time Γ_{Λ_i} enters a peel of Λ_i, its part inside that peel is in one-to-one correspondence with the part of Q bounding that peel minus the regions of P_i and Λ_i that correspond to the external common subtree (because of the projection along segments).

So visiting the same peel of Λ_i twice (necessarily outside these isometric regions of P_i and Λ_i) would produce a contradiction with this bijective correspondence.

We summarize the main import of this section in the following lemma.

Lemma 3 (Cone Embedding) *Let the curve Q on $\mathcal{P}$ live on a cone Λ_i on one side. Then the corresponding half-surface P_i can be isometrically embedded into that cone when it is cut along all edges of C_{P_i} not incident to Q.*

Proof: Because C_{P_i} spans the vertices of P_i, the cutting removes all curvature and leaves a locally flat surface. The peels of P_i may then be embedded within the cone Λ_i via Lemma 1, which guarantees non-overlapping. □ □

6 Source Unfolding

Henceforth let $\overline{P_i}$ be the embedded image of P_i on the cone Λ_i given by Lemma 3. The proof of Theorem 1 below needs one more lemma, concerning cutting and flattening $\overline{P_i}$.

Lemma 4 (Generator Cut) *Let g be a point of Q closest on the cone Λ_i to the apex a of the cone. Then cutting Λ_i along the generator ag unfolds $\overline{P_i}$ to one piece in the plane, i.e., the cut does not disconnect $\overline{P_i}$.*

Proof: We will prove that ag intersects only one peel α_1 of $\overline{P_i}$. In any case, let α_1 be the last peel (from a) of $\overline{P_i}$ intersected by ag. Let y be the closest point to a in $ag \cap \alpha_1$. Then, in α_1, g is the point in Q closest to y, so ag follows a shortest path segment within α_1, projecting to g. By the peel nesting property, this shortest path segment is included in α_1; therefore, cutting along yg doesn't disconnect $\overline{P_i}$.

Suppose now that ag meets another peel α_2 of $\overline{P_i}$ at point x, which projects in α_2 (by a shortest path) to x' on Q. Because $g \in \alpha_1$ and $x' \in \alpha_2$ and α_1, and because α_2 are distinct, g and x' are distinct. Thus $|xx'| < |xg|$, where, to ease notation, we use $|pq|$ to represent $\delta_{\Lambda_i}(p, q)$.

This inequality contradicts the assumption that ag is a shortest path to Q:

$$|ax'| \leq |ax| + |xx'| < |ax| + |xg| = |ag| \ .$$

The first inequality follows from the triangle inequality on Λ_i, which is itself a complete metric space.

Knowing that ag meets just one peel of $\overline{P_i}$ shows that it does not disconnect $\overline{P_i}$, and the lemma claim is established. □ □

We now have assembled all the machinery needed to establish our first main theorem:

Theorem 1 (Half Source Unfolding) *For any Q that lives on a cone Λ_i to the P_i-side, such that each generator of Λ_i meets Q in one point, the source unfolding of the corresponding half P_i of $\mathcal{P}$ is non-overlapping.*

Proof: First, cut all the edges of C_{P_i} not incident to Q. Then apply Lemma 3 to obtain the embedding $\overline{P_i}$ in Λ_i. Finally, cut the generator ag to a closest point $g \in Q$ and unfold $\overline{P_i}$ by Lemma 4 into the plane. □ □

Theorem 1 can be viewed as a significant generalization of the result in [O'R08], which established it for "medial axis polyhedra," whose base edges form Q and whose lateral edges constitute the "upper component" of the cut locus.

The unfolding of the P_1 half of the truncated cube in Figure 4(d) best illustrates this theorem.

We now turn to joining the two halves. Several strategies are available, and we select a simple one. We refer again to Figure 2.

Theorem 2 (Full Source Unfolding) *For Q a convex curve in the class of curves we described, additional cuts permit joining the halves to one, simple polygon.*

Proof: Let P_1 be the half of $\mathcal{P}$ to the convex side of Q, and P_2 the half to the (possibly) nonconvex side. We unfold P_1 just as described in Theorem 1 above. Note that the planar image of Q is a convex curve, as all of Q's vertices are convex to the P_1-side. To the nonconvex side P_2, we cut all of C_{P_2}, including those edges incident to Q. Recall these edges will be incident to vertices that are convex vertices to the P_2-side. In addition, we cut shortest path segments from C_{P_2} to the other (reflex) vertices of Q (none if Q is a quasigeodesic, but possibly several if Q is merely convex to one side). These combined cuts partition P_2 into polygonal regions R_i ($i = 0, 1, 2, \ldots$), each of which projects onto its base $q_{i-1}q_i$ on Q. See Figure 2. The regions are separated by empty cones, whose bounding rays are separated by an angle equal to the curvature at q_i. Because these curvatures are positive, and Q is convex, the joined pieces do not overlap.

It only remains to argue that the cut to g on the P_1-side does not coincide with a cut to a vertex q_i on the P_2-side, for that would disconnect the unfolding into two pieces. This follows from [IIV07, Cor. 1], which shows that g could only be a corner of Q if it were reflex to the P_1-side, because ag makes an angle at least $\pi/2$ with Q to each side of ag. But we know that Q is convex to the P_1-side, so it cannot be that $g = q_i$. □ □

For a convex curve Q shrinking to a point x, the full source unfolding with respect to Q approaches the point source unfolding with respect to x, as one would expect.

7 Future Work

We have not yet addressed the computational complexity of constructing the source unfolding from a given Q, but we expect it will be polynomial in the complexity of Q.

A secondary issue is finding a Q that satisfies our conditions. Although it is known that every convex polyhedron has at least three distinct simple closed quasigeodesics, there is no polynomial-time algorithm known for finding one [DO07, Prob. 24.2]. However, it is easy to find convex curves through at most one vertex: for example, a convex curve inside any face, or the curve obtained by truncating any vertex v of $\mathcal{P}$ orthogonal to a vector within the tangent cone of v (e.g., Figure 3(a)). Also, one can construct quasigeodesic loops through at most one vertex by a minor modification of the technique described in [IOV10]. Such a curve satisfies conditions (a) and (b) in Section 1.

Finally, we leave unresolved determining the largest class of curves Q for which Theorems 1 and 2 hold.

Acknowledgments. We are grateful to the remarkably perceptive comments (and the patience) of Stefan Langerman and several other anonymous referees.

References

[ACC+08] Zachary Abel, David Charlton, Sebastien Collette, Erik D. Demaine, Martin L. Demaine, Stefan Langerman, Joseph O'Rourke, Val Pinciu, and Godfried Toussaint. Cauchy's arm lemma on a growing sphere. Technical Report 90, Smith College, April 2008. arXiv.0804.0986v1 [cs.CG].

[Ale06] Aleksandr D. Alexandrov. *Intrinsic Geometry of Convex Surfaces*. Chapman & Hall/CRC, Boca Raton, FL, 2006. A. D. Alexandrov Selected Works. Edited by S. S. Kutateladze. Translated from the Russian by S. Vakhrameyev.

[AZ67] Aleksandr D. Alexandrov and Victor A. Zalgaller. *Intrinsic Geometry of Surfaces*. American Mathematical Society, Providence, RI, 1967.

[DO07] Erik D. Demaine and Joseph O'Rourke. *Geometric Folding Algorithms: Linkages, Origami, Polyhedra*. Cambridge University Press, July 2007. `http://www.gfalop.org`.

[IIV07] Kouki Ieiri, Jin-ichi Itoh, and Costin Vîlcu. Quasigeodesics and farthest points on convex surfaces. Submitted, 2007.

[IOV09] Jin-ichi Itoh, Joseph O'Rourke, and Costin Vîlcu. Source unfoldings of convex polyhedra with respect to certain closed polygonal curves. In *Proc. 25th European Workshop Comput. Geom.*, pages 61–64. EuroCG, March 2009. Full version submitted to a journal, May 2009.

[IOV10] Jin-ichi Itoh, Joseph O'Rourke, and Costin Vîlcu. Star unfolding convex polyhedra via quasigeodesic loops. *Discrete Comput. Geom.*, 44:35–54, 2010.

[O'R08] Joseph O'Rourke. Edge-unfolding medial axis polyhedra. In *Proc. 24th European Workshop Comput. Geom.*, pages 103–106, March 2008.

[OV11] Joseph O'Rourke and Costin Vîlcu. Conical existence of closed curves on convex polyhedra. `http://arxiv.org/abs/`, February 2011.

[Piz07] Paolo Pizzetti. Confronto fra gli angoli di due triangoli geodetici di eguali lati. *Rend. Circ. Mat. Palermo*, 23:255–264, 1907.

[SS86] Micha Sharir and Amir Schorr. On shortest paths in polyhedral spaces. *SIAM J. Comput.*, 15:193–215, 1986.

[Thu98] W.P. Thurston. Shapes of polyhedra and triangulations of the sphere. *Geometry and Topology Monographs*, 1:511–549, 1998.

Appendix:
Alexandrov-Toponogov Comparison Theorem

Several well-known comparison results for convex surfaces are usually identified as variants of Toponogov's 1959 theorem. For triangles on convex surfaces, however, Alexandrov proved them, without any assumption of differentiability, as early as 1948, and Pizzetti [Piz07] considered the differentiable case in 1907. See [Ale06, p.242] or [AZ67, p.32] for versions in English; see also [ACC$^+$08].

Essentially, the results compare triangles or "hinges" on a given surface to those in the plane.

A *triangle* in a convex surface is a collection of three segments $\gamma_1, \gamma_2, \gamma_3$ such that γ_i and γ_{i+1} have the common endpoint a_{i+2} (indices mod 3). We shall denote the triangle by $\gamma_1\gamma_2\gamma_3$ or, if the segments are clear from the context, by $a_1a_2a_3$. We use $\lambda(\gamma)$ to denote the length of the curve γ.

The first lemma says that if you draw a triangle in the plane with the same lengths as a triangle on a convex surface, the planar triangle angles in general get smaller: they are at most as large as the convex-surface angles.

Lemma 5 *For any triangle $\gamma_1\gamma_2\gamma_3$ in a convex surface there exists a planar triangle $\overline{\gamma}_1\overline{\gamma}_2\overline{\gamma}_3$ with $\lambda(\gamma_i) = \lambda(\overline{\gamma}_i)$. We have $\angle\overline{\gamma}_i\overline{\gamma}_{i+1} \leq \angle\gamma_i\gamma_{i+1}$, $i = 1, 2, 3$ (*mod 3*), and equality holds if and only if $\gamma_1\gamma_2\gamma_3$ is isometric to $\overline{\gamma}_1\overline{\gamma}_2\overline{\gamma}_3$.*

A *hinge* is a pair of segments, γ_1 from a to b and γ_2 from a to c, and the angle $\angle bac$ included between them at a. We denote the hinge by $\gamma_1\gamma_2$.

The second lemma says that if you draw a hinge in the plane with the same angle as a hinge on a convex surface, the planar hinge endpoint separation in general gets larger. Next, we denote by $\overline{\delta}$ the Euclidean distance, and by δ the distance on the surface.

Lemma 6 *For any hinge $\gamma_1\gamma_2$ in a convex surface S there exists a planar hinge $\overline{\gamma}_1\overline{\gamma}_2$ with $\angle\overline{b}\overline{a}\overline{c} = \angle bac$. We have $\delta(b, c) \leq \overline{\delta}(\overline{b}, \overline{c})$, and equality holds if and only if there exists a segment joining b to c in S such that abc is isometric to $\overline{a}\overline{b}\overline{c}$.*