

Star Unfolding Convex Polyhedra via Quasigeodesic Loops*

Jin-ichi Itoh[†] Joseph O’Rourke[‡] Costin Vilcu[§]

July 31, 2009

Abstract

We extend the notion of star unfolding to be based on a quasigeodesic loop Q rather than on a point. This gives a new general method to unfold the surface of any convex polyhedron \mathcal{P} to a simple (non-overlapping) planar polygon: cut along one shortest path from each vertex of \mathcal{P} to Q , and cut all but one segment of Q .

1 Introduction

There are two general methods known to unfold the surface \mathcal{P} of any convex polyhedron to a simple (non-overlapping) polygon in the plane: the source unfolding and the star unfolding. Both unfoldings are with respect to a point $x \in \mathcal{P}$. Here we define a third general method: the star unfolding with respect to a simple closed “quasigeodesic loop” Q on \mathcal{P} . In a companion paper [IOV09], we extend the analysis to the source unfolding with respect to a wider class of curves Q .

The *point source unfolding* cuts the *cut locus* 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 notion of cut locus was introduced by Poincaré [Poi05] in 1905, and since then has gained an important place in global Riemannian geometry; see, e.g., [Kob67] or [Sak96]. The point source unfolding has been studied for polyhedral convex surfaces since [SS86] (where the cut locus is called the “ridge tree”).

* A preliminary version of this work appeared in [IOV07a, IOV07b].

[†]Department of Mathematics, Faculty of Education, Kumamoto University, Kumamoto 860-8555, Japan. j-itoh@kumamoto-u.ac.jp

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

[§]Institute of Mathematics “Simion Stoilow” Romanian Academy, P.O. Box 1-764, RO-014700 Bucharest, Romania. Costin.Vilcu@imar.ro. Partially supported by Romanian Government grant PN II Idei 1187.

The *point star unfolding* cuts the shortest paths from x to every vertex of \mathcal{P} . The idea goes back to Alexandrov [Ale50, p. 181];¹ that it unfolds \mathcal{P} to a simple (non-overlapping) polygon was established in [AO92].

In this paper we extend the star unfolding to be based on a simple closed polygonal curve Q with particular properties, rather than on a single point. This unfolds any convex polyhedron to a simple polygon, answering a question raised in [DO07, p. 307]. The curves Q for which our star unfolding works are quasigeodesic loops, which we now define.

Geodesics & Quasigeodesics. Let Γ be any directed curve on a convex surface \mathcal{P} , and $p \in \Gamma$ be any point in the relative interior of Γ , i.e., not an endpoint. Let $L(p)$ be the total surface angle incident to the left side of p , and $R(p)$ the angle to the right side. Γ is a *geodesic* if $L(p)=R(p) = \pi$ for all such points p . A *quasigeodesic* Γ loosens this condition to $L(p) \leq \pi$ and $R(p) \leq \pi$, again for all p interior to Γ [AZ67, p. 16] [Pog73, p. 28]. So a quasigeodesic Γ has π total face angle incident to each side at all nonvertex points (just like a geodesic), and has at most π angle to each side where Γ passes through a polyhedron vertex. A *simple closed geodesic* is non-self-intersecting (*simple*) closed curve that is a geodesic, and a *simple closed quasigeodesic* is a simple closed curve on \mathcal{P} that is quasigeodesic throughout its length. As all curves we consider must be simple, we will henceforth drop that prefix.

A *geodesic loop* is a closed curve that is geodesic everywhere except possibly at one point, and similarly a *quasigeodesic loop* is quasigeodesic except possibly at one point x , the *loop point*, at which the angle conditions on $L(x)$ and $R(x)$ may be violated—one may be larger than π . Quasigeodesic loops encompass closed geodesics and quasigeodesics, as well as geodesic loops.

Pogorelov showed that any convex polyhedron \mathcal{P} has at least three closed quasigeodesics [Pog49], extending the celebrated earlier result of Lyusternik-Schnirelmann showing the same for differentiable convex surfaces. However, there is no algorithm known that will find a simple closed quasigeodesic in polynomial time: Open Problem 24.2 [DO07, p. 374].

Fortunately it is in general easy to find quasigeodesic loops on a given \mathcal{P} : start at any nonvertex point p , and extend a geodesic from p in opposite directions, following each branch until a self-intersection point is found, either between branches or within one branch. If no vertices are encountered, we have a geodesic loop; if vertices are encountered, at each vertex continue in an arbitrary direction that maintains quasigeodesicity, to obtain a quasigeodesic loop. An exception to this ease of finding a quasigeodesic loop could occur on an *isosceles tetrahedron*: a tetrahedron whose four faces are congruent triangles, or, equivalently, one at which the total face angle incident to each vertex is π . It is proved in [IV08b] that a convex surface possesses a *simple quasigeodesic line*—a non-self-intersecting quasigeodesic infinite (of unbounded length) in both directions—if and only if the surface is an isosceles tetrahedron. So extending a quasigeodesic indefinitely in both directions without self-intersection

¹It is called the “Alexandrov unfolding” in [MP08].

is possible only on isosceles tetrahedra.

Discrete Curvature. The discrete *curvature* $\omega(p)$ at any point $p \in \mathcal{P}$ is the *angle deficit* or *gap*: 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 Gauss-Bonnet theorem says that, for any closed curve Q , the total turn τ of Q (the sum of $\pi - L(p)$ over all p along Q) plus the curvature enclosed to the left of Q , is 2π . For a closed geodesic, $\tau = 0$, so a closed geodesic partitions the curvature into 2π in each “hemisphere” of P . For quasigeodesics that pass through vertices, $\tau \geq 0$ to either side, so the curvature in each half is at most 2π . The curvature in each half defined by a quasigeodesic loop depends on the angle at the loop point.

Some Notation. Denote by P_1 and P_2 the closure of the two components of $P \setminus Q$. As our main focus is usually on one such half, to ease notation we sometimes use P without a subscript to represent either such *half-surface* when the distinction between them does not matter. Unless otherwise stated, vertices of \mathcal{P} are labeled v_i in arbitrary order. We will use pp' to denote a shortest path on \mathcal{P} between p and p' . Other notation will be introduced as needed.

2 Examples and Algorithm

We start with an example. Figure 1(a) shows a geodesic loop Q on the surface \mathcal{P} of a cube. $L(p)=R(p) = \pi$ at every point p of Q except at x , where $R(x)=\frac{3}{2}\pi$ and $L(x)=\frac{1}{2}\pi$. Note that three cube vertices, v_3, v_6, v_7 are to the left of Q , and the other five to the right. This is consistent with the Gauss-Bonnet theorem, because Q has a total turn of $\frac{1}{2}\pi$, so turn plus enclosed curvature is 2π .

For each vertex $v_i \in \mathcal{P}$, we select a shortest path $\text{sp}(v_i) = v_i v'_i$ to Q : a geodesic from v_i to a point $v'_i \in Q$ whose length is minimal among all geodesics to Q . In general there could be several shortest paths from v_i to Q ; we use $\text{sp}(v_i)$ to represent an arbitrarily selected one. The point $v'_i \in Q$ is called a *projection* of v_i onto Q . In this example, for each v_i there is a unique shortest path $\text{sp}(v_i)$, which is the generic situation.

Algorithm. If we view the star unfolding as an algorithm with inputs \mathcal{P} and Q , it consists of four main steps:

0. Cut \mathcal{P} along Q , producing the half-surfaces P_1 and P_2 .
1. For each half-surface P_k , select shortest paths $\text{sp}(v_i)$ from each $v_i \in P_k$ to Q .
2. Cut along all $\text{sp}(v_i)$ and flatten each half-surface.
3. Rejoin the two flattened halves at a segment $s \subset Q$.




Figure 1: (a) Geodesic loop Q on cube. Shortest paths $sp(v_i)$ are shown. Faces are labeled $\{F, T, L, R, Bt, Bk\}$. (b) Star unfolding with respect to Q , joined at $s = v'_0v'_7$.

After cutting along $sp(v_i)$, we conceptually insert an isosceles triangle with apex angle $\omega(v_i)$ at each v_i , which flattens each resulting half. One half-surface (in our example, the left half), is convex, while the other resulting half-surface may have several points of nonconvexity, at the images of x . (In our example, only the image x_1 is nonconvex, when the inserted “curvature triangles” are included.) In the final step of the procedure, we select a segment s of Q whose interior contains neither a vertex v_i nor any vertex projection v'_i , such that the extension of s is a supporting line of each half, and glue to the halves along s . In our first example, we choose $s = v'_0v'_7$ (many choices for s work in this example), which leads to non-overlap of the two halves.

We illustrate with one more example before proceeding. Again \mathcal{P} is a cube, but now Q is the closed quasigeodesic composed of the edges bounding one square face, the bottom face in Figure 2(a). Cutting four shortest paths from the other vertices orthogonal to Q , and cutting all but one edge s of Q , results in the standard Latin-cross unfolding of the cube shown in (b).

We now proceed to detail the three steps of the procedure, this time with proofs. Because the proofs for quasigeodesics are straightforward in comparison to the proofs for quasigeodesic loops, we separate the two in the exposition.

3 Quasigeodesics

3.1 Shortest Path Cuts for Quasigeodesics

We again use a cube as an illustrative example, but this time with a closed quasigeodesic Q , not a loop: $Q = (v_0, v_5, v_7)$; see Figure 3(a). There is π angle incident to the right at v_5 , and $\frac{1}{2}\pi$ incident to the left; and similarly at v_0 and v_7 . At all other points $p \in Q$, $L(p) = R(p) = \pi$. Thus Q is indeed a quasigeodesic. We will call the right half (including v_2) P_1 , and the left half (including v_4) P_2 .

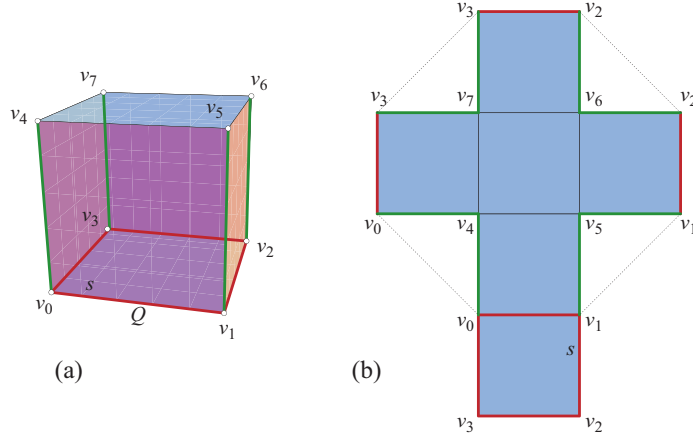


Figure 2: (a) $Q = (v_0, v_1, v_2, v_3)$. (b) Star unfolding with respect to Q . Here the uncut segment of Q is $s = v_0v_1$.

In Figure 3(a), the paths from $\{v_1, v_3, v_6\}$ are uniquely shortest. From v_2 there are three paths tied for shortest, and from v_4 also three are tied.

A central fact that enables our construction is this key lemma from [HIV07, Cor. 1], slightly modified for our circumstances:

Lemma 1 *Let W be a simple closed polygonal curve on a convex surface \mathcal{P} , and let p be any point of one (closed) half-surface P bounded by W , but not on W . Let p' be one of the points of W closest to p . Then for any choice of $\text{sp}(p) = pp'$, the angle α made by $\text{sp}(p)$ with W at p' is at least $\pi/2$. In particular, if p' is not a corner of W then $\alpha = \pi/2$ and the path $\text{sp}(p)$ is unique as shortest between p and p' , and $\alpha > \pi/2$ occurs only at corners p' where the angle of W toward P is larger than π .*

A second fact we need concerning the shortest paths $\text{sp}(v_i)$ is that they are disjoint.

Lemma 2 *Any two shortest paths $\text{sp}(v_1)$ and $\text{sp}(v_2)$ are disjoint, for distinct vertices $v_1, v_2 \in P$.*

Proof: Suppose for contradiction that at least one point u is shared: $u \in \text{sp}(v_1) \cap \text{sp}(v_2)$. We consider four cases: one shortest path is a subset of the other, the shortest paths cross, the shortest paths touch at an interior point but do not cross, or their endpoints coincide.

1. $\text{sp}(v_2) \subset \text{sp}(v_1)$. Then $\text{sp}(v_1)$ contains a vertex v_2 in its interior, which violates a property of shortest paths [SS86, Lem. 4.1].
2. $\text{sp}(v_1)$ and $\text{sp}(v_2)$ cross properly at u . It must be that $|uv'_1| = |uv'_2|$, otherwise both paths would follow whichever tail is shorter. But now it




Figure 3: (a) Cube and quasigeodesic $Q = (v_0, v_5, v_7)$. Shortest paths $sp(v_i)$ as indicated. (b) Flattening the right half-surface by insertion of curvature triangles along the shortest paths $sp(v_i) = v_i v'_i$. (c) Flattening the left half-surface. (d) Two halves joined at $s = v_5 v'_6$.

is possible to shortcut the path in the vicinity of u via σ as shown in Figure 4(a), and the path (v_1, σ, v'_2) is shorter than $\text{sp}(v_1)$.

3. $\text{sp}(v_1)$ and $\text{sp}(v_2)$ touch at u but do not cross properly there. Then there is a shortcut σ to one side (the side with angle less than π), as shown in Figure 4(b).
4. $v'_1 = v'_2$. From Lemma 1 we know the two paths are orthogonal to the quasigeodesic Q at $v'_1 = v'_2$ hence, since they cannot be in the situation of Case 1, there is an angle $\alpha > 0$ separating the paths in a neighborhood of the common endpoint; see Figure 4(c). Then Q has more than π angle to one side at this point, violating the definition of a quasigeodesic.

□




Figure 4: Lemma 2: (a) paths cross; (b) paths touch at an interior point; (c) paths meet at endpoint.

This lemma ensures that the cuttings along $\text{sp}(v_i)$ do not interfere with one another.

3.2 Flattening the Half-Surfaces determined by Quasigeodesics

The next step is to flatten each chosen half-surface P_1 and P_2 (independently) by suturing in “curvature triangles” along each $\text{sp}(v)$ path. Let P be one of P_1 or P_2 . The basic idea goes back to Alexandrov [Ale05, p. 241, Fig. 103], and was used also in [IV08a]. Let ℓ be the length of a shortest path $\text{sp}(v)$, and let $\omega = \omega(v) > 0$ be the curvature at v . We glue into $\text{sp}(v) = vv'$ the isosceles *curvature triangle* Δ with apex angle ω gluing to v , and incident sides of length ℓ gluing along the cut vv' . This is illustrated in Figure 3(b,c). We display this in the plane for convenience of presentation; the triangle insertion should be viewed as operations on the manifolds P_1 and P_2 , each independently.

This procedure only works if $\omega < \pi$, for ω becomes the apex of the inserted triangle \triangle . If $\omega \geq \pi$, we glue in two triangles of apex angle $\omega/2$, both with their apexes at v .² Slightly abusing notation, we use \triangle to represent these two triangles together. In fact we must have $\omega < 2\pi$ for any vertex v (else there would be no face angle at v), so $\omega/2 < \pi$ and this insertion is indeed well defined.

We should remark that an alternative method of handling $\omega \geq \pi$ would be to simply not glue in anything to the vertex v with $\omega(v) \geq \pi$, in which case we still obtain the lemma below leading to the exact same unfolding.

Now, because ω is the curvature (angle deficit) at v , gluing in \triangle there flattens v to have total incident angle 2π . Thus v is no longer a vertex of P (and two new vertices are created along the bounding curve).

Let P^Δ be the new manifold with boundary obtained after insertion of all curvature triangles into P . We claim that a planar embedding of P^Δ does not overlap; i.e., P^Δ is isometric to a simple planar polygonal domain.

In the proof we use two results of Alexandrov. The first is his celebrated theorem [Ale05, 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.

The second is a tool that relies on this theorem:

Lemma 3 *Let M be a convex polyhedral manifold with convex boundary, i.e., the angle toward M of the left and right tangent directions to ∂M is at most π at every point. Then the closed manifold $M^\#$ obtained by gluing two copies of M along ∂M by identifying the corresponding points is isometric to a unique polyhedron, with a plane of symmetry Π containing ∂M , which is a convex polygon in Π .*

Proof: That $M^\#$ is a convex polyhedron follows from Alexandrov's gluing theorem. Because $M^\#$ has intrinsic symmetry, a lemma of Alexandrov [Ale05, p. 214] applies to show that the polyhedron has a symmetry plane containing the polygon ∂M . \square

One consequence of this lemma we record in a corollary:

Corollary 4 *A convex polyhedral manifold M with convex boundary Q and with no interior curvature is isometric to a planar convex polygon.*

Proof: Apply Lemma 3 to M and let $M^\#$ be the resulting polyhedron. The lemma says that $M^\#$ has a symmetry plane containing $Q = \partial M$. As all the vertices of $M^\#$ are on Q , $M^\#$ itself must be planar. Therefore Q is the boundary of a planar convex polygon. \square

Now we apply these tools to P^Δ :

Lemma 5 *For each half-surface P of \mathcal{P} , P^Δ is a planar convex polygon, and therefore simple (non-overlapping).*

²One can view this as having two vertices with half the curvature collocated at v .

Proof: P^Δ is clearly a topological disk: P is, and the insertions of Δ 's maintains it a disk. At every interior point of P^Δ , the curvature is zero by construction. So the interior is flat.

Next we show that the boundary ∂P^Δ is convex. This follows from the orthogonality of $\text{sp}(v)$ guaranteed by Lemma 1, as the base angle of the inserted triangle(s) is $\pi/2 - \omega/2$ for $\omega < \pi$, or $\pi/2 - \omega/4$ for $\omega \geq \pi$ (see Figure 5; $\omega = \omega(v)$), so the new angle is smaller than π by $\omega/2$ or $\omega/4$.




Figure 5: Lemma 5, Case 1: (a) $\omega < \pi$; (b) $\omega \geq \pi$.

Once we know that P^Δ is homeomorphic to a disk and has a convex boundary, Corollary 4 shows that P^Δ is a planar convex polygon, and therefore simple. \square

Note that, when the total curvature in P is 2π then the straight development of Q is turned 2π by the Δ insertions, as in Figure 3(b). When the total curvature in P is less than 2π , the development of Q is not straight, but the Δ insertions turn it exactly the additional amount needed to close it to 2π , as in (c) of the same figure.

3.3 Joining the Halves for Quasigeodesics

The final step of the unfolding procedure selects a *supporting segment* $s \subset Q$ whose relative interior does not contain a projection v' of a vertex. All of Q will be cut except for s . When Q is a closed quasigeodesic, any choice for s generates a supporting line to a planar development of P_k , $k = 1, 2$, because P_k^Δ is a convex domain. Then joining planar developments of P_1 and P_2 along s places them on opposite sides of the line through s , thus guaranteeing non-overlap. See Figure 3(d), where $s = v_5 v'_6$.

4 Quasigeodesic Loops

We now turn to the case when Q is a closed quasigeodesic loop with loop point x , at which the angle toward P is $\beta > \pi$. See Figure 1, where $\beta = \frac{3}{2}\pi$.

4.1 Shortest Path Cuts for Quasigeodesic Loops

Lemma 1, claiming that the shortest path $\text{sp}(p) = pp'$ is unique and orthogonal to Q at p' , needs no modification. For $p' = x$, the shortest path px makes a non-acute (at least $\pi/2$) angle with Q on both sides. Lemma 2, claiming disjointness of the shortest paths for distinct vertices v_1, v_2 , need only be modified by observing that more than one shortest path can project to x , but the paths do not intersect other than at x . The proof is identical.

The significant differences between quasigeodesics and quasigeodesic loops are concentrated largely in the proofs for flattening and for joining the halves. We should emphasize that the algorithm steps detailed earlier are the same; it is only the justification that becomes more complicated.

4.2 Flattening the Half-Surfaces determined by Quasigeodesic Loops

We aim to prove the analog of Lemma 5 for quasigeodesic loops: P^Δ is (isometric to) a simple, planar polygon. Obviously that lemma establishes this for the convex half-surface of \mathcal{P} . In this subsection let P be the nonconvex half-surface, in which the angle β at x exceeds π .

The argument will be different (and easier) if no vertices project to x .

4.2.1 No Vertex Projects to x

Lemma 6 *If no vertex of P projects to x , then P^Δ is a simple, planar polygon.*

Proof: Let γ be a geodesic on P^Δ bisecting the angle β at x , and meeting ∂P^Δ at y ; see Figure 6. Because P^Δ is flat, γ does in fact meet the boundary of P^Δ , i.e., it does not self-cross before reaching the boundary. Cut P^Δ along $\gamma = xy$ into two manifolds. Because $\beta < 2\pi$, the angle at x in each manifold is less than π . Therefore, we have obtained exactly the situation in Lemma 5: each is a flat manifold homeomorphic to a disk, with a convex boundary. Corollary 4 then implies that each is isometric to a planar convex polygon, with a supporting line through xy . Therefore they may rejoin along γ to form a planar simple polygon, with one reflex (concave) vertex at x . \square

4.2.2 Vertices project to x

Now we consider the more difficult case, when vertices in a set $V = \{v_1, v_2, \dots, v_k\}$ project to x . Rather than construct P^Δ directly as above, we proceed in two stages: first we insert curvature triangles for all those vertices of P that project to $Q \setminus \{x\}$, forming a (nonflat) manifold X , and then argue that X unfolds without overlap. We will not insert curvature triangles for the vertices in V , but rather argue directly for non-overlap of a flattening $U(X)$ of X . The boundary of $U(X)$ will consist of two portions: a convex chain Q^c deriving from $Q \setminus \{x\}$




Figure 6: No vertex of P projects onto Q at x .

and the insertion of the curvature triangles, and the remainder, a (generally) nonconvex chain Q^n deriving from x and the vertices in V . Much as in Figure 6, we will cut X by a geodesic γ and later glue the two unfolding pieces back along γ . We will see that the chain Q^n is formed of subchains of particular star unfoldings with respect to x .

Preliminaries. We need several ingredients for this proof. First, we will need the cut locus of x on the manifold-with-boundary X , which may be defined as in Section 1: $C(x)$ is the closure of the set of all points of X joined to the source point x by at least two shortest paths on X .

Second, we need two extensions to the theorem established in [AO92, Thm. 9.1] that the star unfolding of a polyhedron does not overlap. That paper assumed the source point x was generic in the sense that it had a unique path to every vertex. We extend the result to the case when x has two or more distinct shortest paths to a vertex v : then simply cutting one of the paths to v again leads to non-overlap. The second extension is more substantive. For any three points r, s, t in the plane, forming the counterclockwise angle $\theta = \angle rst$ at s , define the *angular region* determined by these points to be the set of points p such that the segment ps falls within the angle θ : counterclockwise of rs and clockwise of ts .

Lemma 7 *Let v be a vertex of the star unfolding $S(x)$ of a polyhedral convex surface \mathcal{P} with respect to $x \in \mathcal{P}$, and x_1 and x_2 the images of x resulting from cutting a shortest path xv , with x_2 counterclockwise of x_1 around $S(x)$. Then $S(x)$ is disjoint from the open angular region determined by x_1, v, x_2 , or, equivalently, $S(x)$ is enclosed in the closed angular region determined by x_2, v, x_1 .*

See Figure 7. Both of these extensions to [AO92] are established in [OV09].

We will also need this fact:




Figure 7: The open angular region determined by x and the vertex v is disjoint from the star unfolding $S(x)$.

Lemma 8 *Let $M^\#$ be a convex polyhedron obtained by doubling M as in Lemma 3. Then a shortest path vu from a vertex v on the symmetry plane Π , to a vertex u not on Π , lies completely in the half containing u .*

Proof: Let u be in the bottom half. Suppose for contradiction that a shortest path vu includes a section σ that enters the top half at $s_1 \in \Pi$ and exits back to the bottom half at $s_2 \in \Pi$. Then we can reflect σ through Π to place it on the bottom half, without changing its length. Because shortest paths on convex surfaces do not branch, we obtain a contradiction. \square

The flattening proof consists of two parts: first, a procedure is used to find the splitting geodesic γ , or to determine no such geodesic exists. The second part then has two cases for reaching the final unfolding.

Part I: Excising Digons. We now describe a process to flatten a subset of X by cutting out regions containing all its interior vertices. The flattened version X' will then be used in further steps of the proof.

Recall we defined X as the manifold obtained by inserting curvature triangles into P for all shortest paths incident to $Q \setminus \{x\}$, but leaving $V = \{v_1, v_2, \dots, v_k\}$ unaltered. So X is not flat, but its boundary ∂X is the simple closed convex curve Q^c .

Let $C(x)$ be the cut locus of x on X . In general $C(x)$ is a forest of trees, each of which meets ∂X in a single point, and whose leaves are the vertices in

V. This follows from the proof of Theorem A in [ST96], but may also be seen intuitively as follows. By a remark of Alexandrov [Ale05, p. 236], there exist many convex polyhedral surfaces \mathcal{X} containing X . The cut locus of x on \mathcal{X} is a tree, which gets clipped by ∂X to a forest on X , provided \mathcal{X} can be chosen such that any shortest path in ∂X starting at x crosses ∂X only once. A direct way to see this is to observe that, at every vertex v interior to X , an edge of $C(x)$ emanates, and when two such edges meet interior to X , they join to start a third edge of $C(x)$.

Let w_1, w_2, \dots be the points of $C(x) \cap \partial X$. Each w_i is joined to x by at least two shortest paths, and the union of two of these bounds a digon D_i that includes the tree component T_i of $C(x)$ incident to w_i .

Cut out from X all D_i 's, and glue back their boundary segments in $X \setminus \bigcup_i \text{int} D_i$ to get a new surface X' . Because the vertices V of X have been excised with the digons, X' is flat. And $\partial X'$ is convex except (possibly) at x . Therefore, the argument of Lemma 6 applies to X' , showing that it is the join of two planar convex polygons.

We illustrate this construction before proceeding with the proof, first with a real example, second with an abstract example. Figure 8(a) illustrates X for the geodesic loop on a cube from Figure 1; only v_1 projects to x . $C(x)$ is in this case the single segment xw_1 . After excising the digon bounded by the two xw_1 shortest paths, the planar polygon X' shown in (b) is attained.




Figure 8: (a) Manifold X for the nonconvex half-surface of Figure 1(b). The two points labeled x should be identified along the shortest path v_1x ; so X in this case is isometric to a subset of a cone with apex v_1 . (b) After removing the digon bounded by paths from x to w_1 and identifying the paths.

A more generic situation is illustrated in Figure 9.




Figure 9: (a) The manifold X depicted abstractly. (b) Planar, simple X' after excision of digons.

Part II: Partitioning X , Case 1. The next step of the proof is to partition the manifold X into pieces each of which has a convex boundary. The partition will be into either two or three pieces. We start with the two-piece Case 1. Let α_i be the digon angle of D_i at x on X . Then the angle β at x on X is reduced by $\sum \alpha_i$ on X' . Case 1 holds when there is a segment $\gamma = xy$ on X' from x to some point $y \in \partial X'$ such that the corresponding geodesic on X splits the angle β at x into two parts, each of which is at most π . (It is not possible to simply select the bisector of β on X for γ , because that bisector might fall inside the α_i of one of the digons, and so not be realized in X' .) Assuming for Case 1 that such a γ exists, we partition X by cutting along γ , resulting in two (nonflat) manifolds X_1 and X_2 . Note that each such manifold has a convex boundary: ∂X is convex at every point excepting x , the angle at y is less than π , and we have conveniently split the angle at x . Therefore, we may apply the “doubling” Lemma 3 to each half. Henceforth we confine our attention to X_1 .

According to that lemma, we have a convex polyhedron $X_1^\#$ with ∂X_1 lying in the symmetry plane Π . We’ll call the portion above and below Π the upper and lower halves respectively. The upper half includes (a subset of) vertices of V , say v_1, \dots, v_j , and the lower half includes equivalent copies, call them u_1, \dots, u_j . There are a number of vertices on Π deriving from ∂X_1 , including x, y ; call them z_0, z_1, z_2, \dots , with z_0 adjacent to y .

Now the plan is to construct the star unfolding of $X_1^\#$ with respect to the point x . By Lemma 8, the shortest paths from x to the vertices of $X_1^\#$ in Π lie wholly in one half or the other. For a vertex on Π , either the shortest path lies in Π , or there are pairs of equal-length shortest paths, one in the upper and one in the lower half. By our extension of the star unfolding theorem, we may choose which shortest path to cut in the case of ties, and still obtain a non-overlapping unfolding. We choose to select all shortest paths to the vertices on Π in the lower half (or along Π if that is where they lie). See Figure 10(a). Note that $\gamma = xy$ will necessarily lie in Π , and it is also necessarily a shortest path from x

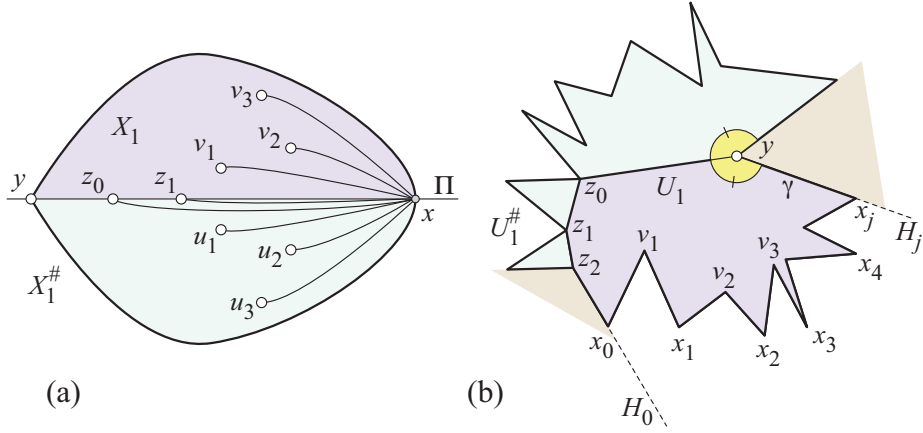


Figure 10: (a) The double polyhedron $X_1^\#$ of X_1 . (b) The star unfolding $U_1^\#$ of $X_1^\#$ with respect to x .

to y (because it is an edge of the polyhedron, and every edge of a polyhedron is a shortest path between its endpoints). So we cut $xy = \gamma$ as well. The result is a planar, non-overlapping unfolding $U_1^\#$. Now, identify within $U_1^\#$ the portion U_1 corresponding to the upper half. Also note that γ is one edge of ∂U_1 . Let x_{i-1} and x_i be the images of x adjacent to v_i . See Figure 10(b). Note that $Q_1^n = (x_0, v_1, x_1, v_2, \dots, v_j, x_j) \subset \partial U_1$ is a subchain of the nonconvex chain Q^n mentioned earlier, and its complementary chain Q_1^c in ∂U_1 is a subchain of the convex chain Q^c .

Next we perform the exact same procedure for X_2 , resulting in $U_2^\#$ containing U_2 . Now glue $U_1^\#$ to $U_2^\#$ along their boundary edges deriving from γ . Applying Lemma 7 at the point y in each unfolding shows that they join without overlap, as follows. Let θ_1 be the angle at y in U_1 , and θ_2 the angle in U_2 . If y coincides




Figure 11: Complementary exterior open angular regions at y .

with a vertex of Q , then $\theta_1 + \theta_2 < \pi$, hence joining the closed angular regions enclosing $U_1^\#$ and $U_2^\#$ leaves at y an empty open angular region of measure

$2(\pi - \theta_1 - \theta_2)$. If y is not a vertex of Q , then $\theta_1 + \theta_2 = \pi$, hence the closed angular regions enclosing $U_1^\#$ and $U_2^\#$ are complementary; see Figure 11. Thus, in either case, the contained U_1 and U_2 do not overlap one another. And so we have established that $U(X) = U_1 \cup U_2$ is a planar simple polygon.

Part II: Partitioning X , Case 2. Case 1 relies on the existence of a segment γ on X' such that the corresponding geodesic on X splits the angle β at x into two parts, each of which is at most π . Now we consider the possibility that there is no such γ . We illustrate this possibility with an example before handling this case.

In Figure 12, \mathcal{P} is a tetrahedron with three right angles incident to v_0 . The geodesic loop shown has at x an angle $\beta = 330^\circ$ in the lower half, and two vertices v_1 and v_2 , included in the same digon, projecting to x . In this example, the sole digon constitutes the majority of X , and its removal leaves X' ((c) of the figure) so narrow as to not admit a γ with the desired angle-splitting properties.

In this case, we have a “fat” digon D_i (possibly $D_i = X$) whose angle α_i at x covers all the possible splitting segments γ . Let $D = D_i$ to ease notation, and let γ_1 and γ_2 be the two boundary edges (shortest paths) of D connecting x to w_i . Let a and b be the vertices adjacent to x on ∂X . The angle from ax to γ_1 , and the angle from bx to γ_2 , are both less than π . Now we define three (topologically) closed submanifolds of X : X_1 , the portion bounded by γ_1 and containing vertex a , X_2 , the portion bounded by γ_2 and containing vertex b , and the digon D in between. X_1 and X_2 have convex boundaries, just as the manifolds in Case 1, and we go through the identical process:

$$X_i \rightarrow X_i^\# \rightarrow U_i^\# \rightarrow U_i.$$

The digon D might not be convex at x ; in Figure 12(b) the angle at x is $\alpha_i = 300^\circ$. So we cannot use the doubling lemma. Instead we glue (“zip”) γ_1 to γ_2 , which produces a convex polyhedron D^z containing the vertices of V inside D , and vertices at x and at w_i . Now we produce the star unfolding of D^z with respect to x . Note that $xw_i = \gamma_1 = \gamma_2$ is a shortest path on D^z (because γ_1 and γ_2 are shortest paths on $X \supseteq D$, there can be no shorter path in D , and a path crossing $\gamma_1 = \gamma_2$ on D^z is necessarily longer), so that gets cut as part of the star unfolding. Call this unfolding U_D . By Lemma 7, U_D is contained within the closed angular region at w_i , bounded by rays along γ_1 and γ_2 .

Next we join $U_1^\#$ to U_D along γ_1 , and $U_2^\#$ to U_D along γ_2 , non-overlapping by the angular region properties, and conclude that the contained $U(X) = U_1 \cup U_D \cup U_2$ is a planar simple polygon. See Figure 12(d).

Finally, removing the inserted curvature triangles from $U(X)$ establishes the analog of Lemma 5:

Lemma 9 *For Q a quasigeodesic loop, the star unfolding of the nonconvex half-surface P of \mathcal{P} is a planar, simple polygon.*

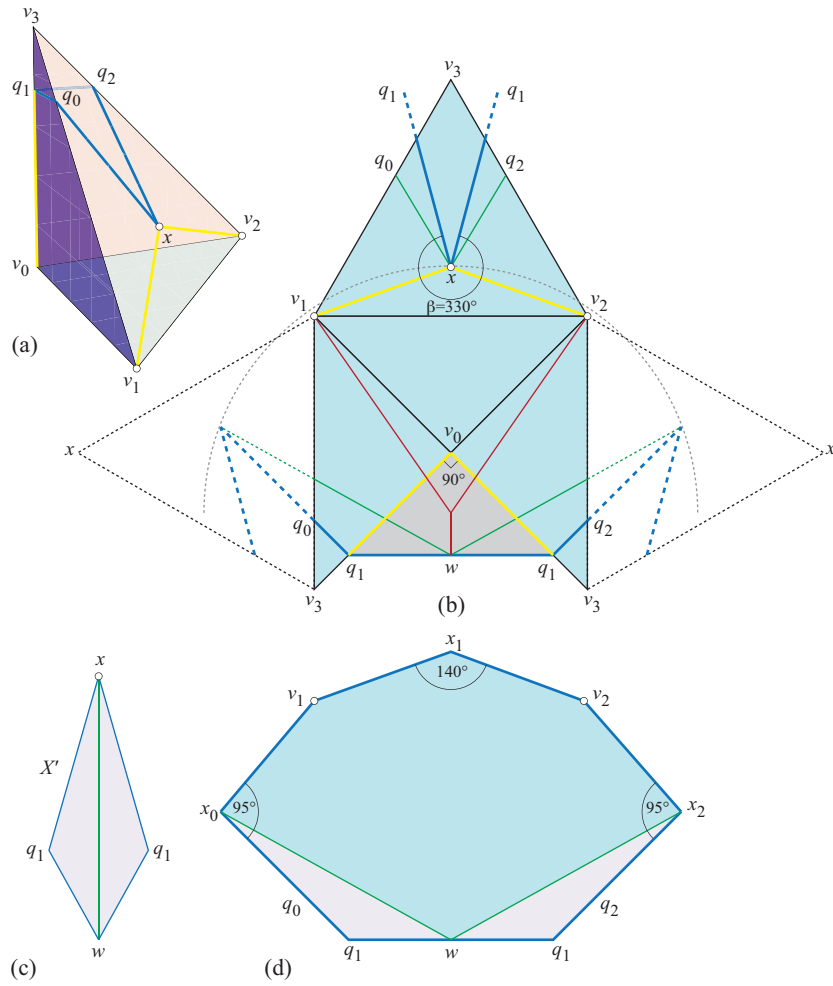


Figure 12: (a) Geodesic loop $Q = (q_0, q_1, q_2, x)$. (b) The manifold after insertion of the curvature triangle with $\omega(v_0) = 90^\circ$ at v_0 . $C(x)$ is the ‘Y’, and the digon is bounded by xw . (c) X' : after removal of digon. (d) $U(X)$. Note the angles at the three images of x sum to $\beta = 330^\circ$.

4.3 Joining the Halves for Quasigeodesic Loops

Let s be the edge of the unfolding $U = U(X)$ on which y lies (in Case 1) or on which w lies (in Case 2). If y or w is at a vertex of \mathcal{P} , then let s be either incident edge. We claim that the line L containing s is a supporting line to $U(X)$. We establish this by identifying a larger class of supporting lines, which includes L . We make the argument for Part II, Case 1 above (when γ exists), as Case 2 is very similar.

Returning to Figure 10(b), consider the two empty angular regions exterior to $U_1^\#$, incident to the vertex adjacent to x_0 that is not v_1 (z_2 in that figure), and the vertex adjacent to x_j that is not v_j , i.e., y . Let H_0 and H_j be the halflines on those angular region boundaries that include x_0 and x_j respectively. The empty angular regions imply that U_1 is included in the convex region \bar{U}_1 of the plane delimited by H_0 , H_j , and the convex boundary $Q_1^c \subset \partial U_1$: the nonconvex boundary portion $Q_1^n \subset \partial U_1$ can cross neither H_0 nor H_j . Let p_1 be the point in $H_0 \cap H_j$, which might be “at infinity” if those halflines diverge. Similarly, U_2 is included in a region \bar{U}_2 delimited by H_j and H_k , which meet at p_2 or not at all if those halflines diverge. Let L_j be the line containing H_j , which contains γ , and p_1 and p_2 when those exist. We consider three cases.

1. *Both p_1 and p_2 exist.* See Figure 13(a). Then there are two supporting lines parallel to L_j , and every edge between them along ∂U is supporting to \bar{U} and so to U .
2. *p_2 exists, but p_1 does not.* See Figure 13(b). We have the one line parallel to L_j supporting \bar{U}_2 . Let p' be a point on L_j on the opposite side of y from p_2 that maximizes the chain of $\partial \bar{U}_1$ visible from p' , say the point that minimizes $|p'y|$. Then again every edge between the two supporting lines—that parallel to L_j supporting \bar{U}_2 , and the line through p' supporting \bar{U}_1 —is supporting to \bar{U} and so to U .
3. *Neither p_1 nor p_2 exist.* See Figure 13(c). Then again let p' be a point on L_j on the opposite side of y that maximizes the visibility of $\partial \bar{U}$. Then all the edges between the tangency points extend to supporting lines.

In particular, we see that the line(s) extending edge(s) incident to y are among the supporting lines, as claimed. For example, s is $v'_0 v'_4$ in Figure 8(b), and that edge is used to join the halves in Figure 1(b).

One issue remains. It could be that the segment s identified above is the base of a curvature triangle, rather than a segment of Q , in which case it cannot be used for joining the halves. Returning to Figure 5, there are two cases: (a) $\omega < \pi$, and (b) $\omega \geq \pi$. In case (a), s is the base of a curvature triangle, with the angle at either endpoint of the base larger than $\pi/2$. Note that edges adjacent to s must be edges of Q . Regardless of the angle at which L_j intersects s , one of these two adjacent edges must include supporting portion of ∂U , for example, visible from p' in Figure 13(b,c). In case (b), two curvature triangles are inserted, but as we noted earlier, neither is truly needed, for the boundary




Figure 13: Lines determined by the marked portion of ∂U are supporting for $U = U_1 \cup U_2$.

of ∂U (∂P in Lemma 5) is already convex at the apexes of the two curvature triangles. So, simply not inserting them leaves an edge of Q crossed by L_j that can serve as s . So, in all situations, we obtain a supporting segment.

5 Conclusion

We have established our main theorem:

Theorem 1 *Let Q be a quasigeodesic loop on a convex polyhedral surface \mathcal{P} . Cutting shortest paths from every vertex to Q , and cutting all but a supporting segment s of Q as designated above, unfolds \mathcal{P} to a simple planar polygon.*

Figure 14 shows another example, a closed geodesic on a dodecahedron, this time a pure geodesic. The unfolding following the above construction is shown in Figure 14(c,d). In this case when Q is a pure, closed geodesic, there is additional structure that can be used for an alternative unfolding. For now Q lives on a region isometric to a right circular cylinder. Figure 14(b) illustrates that the upper and lower rims of the cylinder are loops parallel to Q through the vertices of P at minimum distance to Q (at least one vertex on each side.) In the figure, these shortest distances to the upper rim are the short vertical paths from Q to the five pentagon vertices. Those rim loops are themselves closed quasigeodesics. An alternative unfolding keeps the cylinder between the rim loops intact and attaches the two reduced halves to either side. See Figure 15.

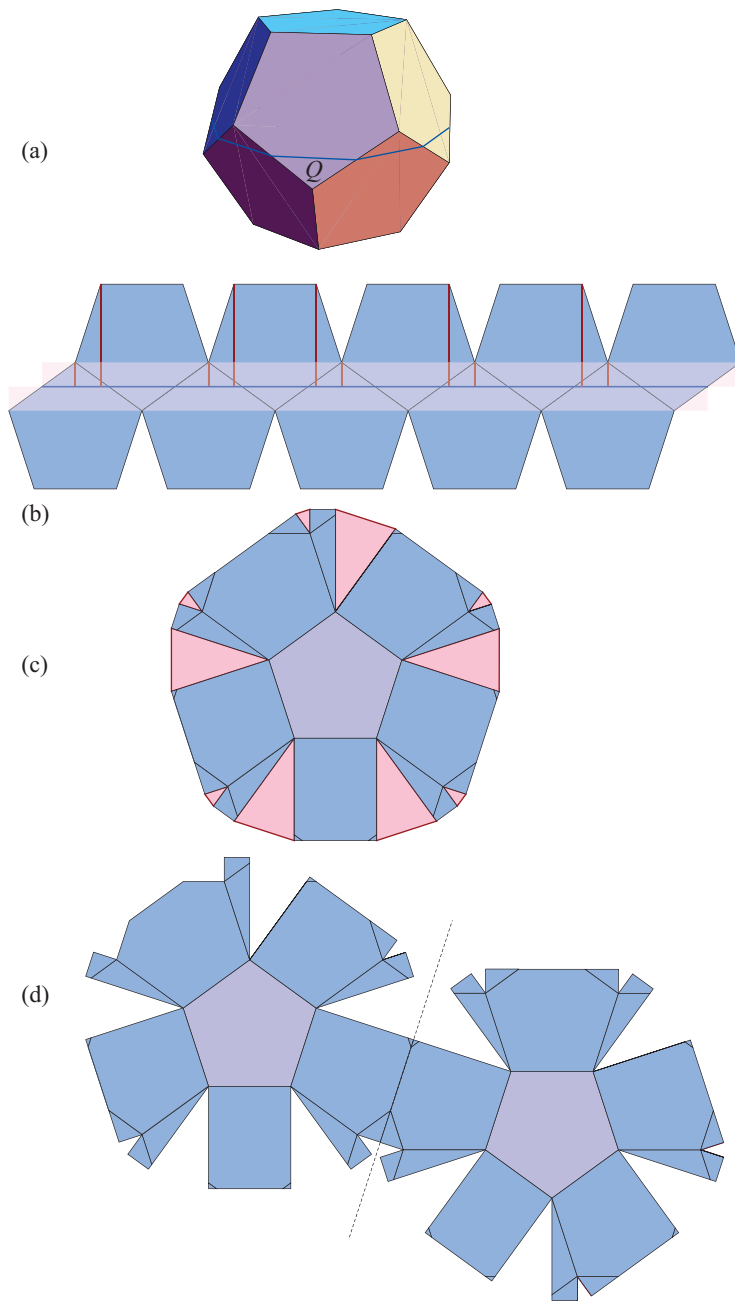


Figure 14: (a) Q here is a geodesic; it includes no vertices, as is evident in the layout (b). The region isometric to a right circular cylinder is highlighted. The convex domain P^Δ from Lemma 5 is shown in (c), and one possible unfolding in (d).

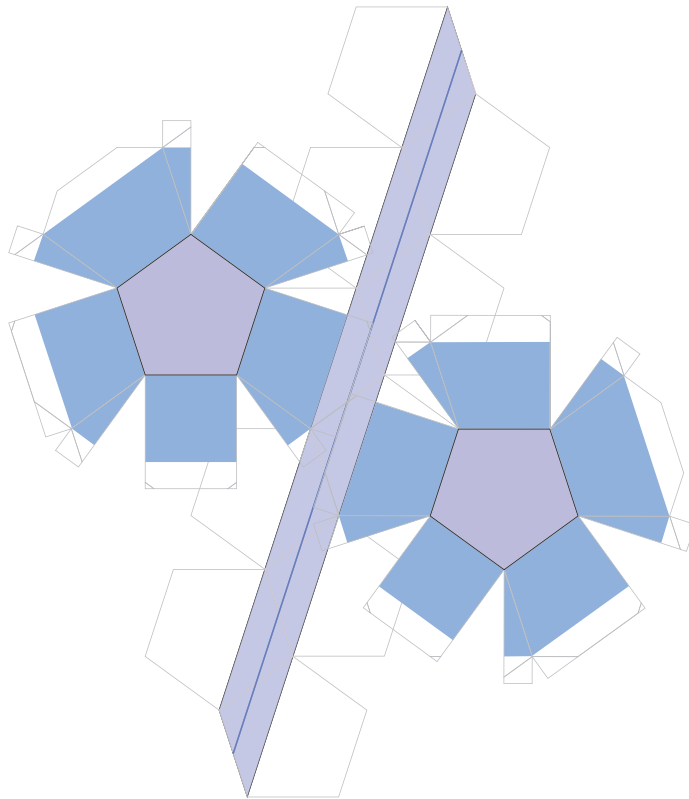


Figure 15: Alternative unfolding of the example in Figure 14. Various construction lines are shaded lightly.

5.1 Future Work

We have focused on establishing Theorem 1 rather than the algorithmic aspects. Here we sketch preliminary thoughts on computational complexity. Let n be the number of vertices of \mathcal{P} , and let $q = |Q|$ be the number of faces crossed by the geodesic loop Q . In general q cannot be bounded as a function of n . Let $m = n + q$ be the total combinatorial complexity of the “input” to the algorithm. Constructing Q from a given point and direction will take $O(q)$ time. Identifying a supporting segment s , and laying out the final unfolding, is proportional to m . The most interesting algorithmic challenge is to find the shortest paths from each vertex v_i to Q . The recent $O(n \log n)$ algorithm in [SS08] leads us to expect the computation can be accomplished efficiently.

We do not believe that quasigeodesic loops constitutes the widest class of curves for which the star unfolding leads to non-overlap. Extending Theorem 1 to quasigeodesics with two exceptional points, one with angle larger than π to one side, and the other with angle larger than π to the other side, is a natural next step, not yet completed.

If one fixes a nonvertex point $p \in \mathcal{P}$ and a surface direction \vec{u} at p , a quasigeodesic loop can be generated to have direction \vec{u} at p . It might be interesting to study the continuum of star unfoldings generated by spinning \vec{u} around p .

Acknowledgments. We thank Boris Aronov for many observations and suggestions which improved the paper.

References

- [Ale50] Aleksandr D. Alexandrov. *Vypyklyue Mnogogranniki*. Gosydarstvennoe Izdatelstvo Tehno-Teoreticheskoi Literatury, 1950. In Russian. See [Ale58] for German translation, and [Ale05] for English translation.
- [Ale58] Aleksandr D. Alexandrov. *Konvexe Polyeder*. Akademie Verlag, Berlin, 1958. Math. Lehrbücher und Monographien. Translation of the 1950 Russian edition.
- [Ale05] 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.
- [AO92] Boris Aronov and Joseph O’Rourke. Nonoverlap of the star unfolding. *Discrete Comput. Geom.*, 8:219–250, 1992.
- [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 Vilcu. Quasigeodesics and farthest points on convex surfaces. Submitted, 2007.
- [IOV07a] Jin-ichi Itoh, Joseph O’Rourke, and Costin Vilcu. Unfolding convex polyhedra via quasigeodesics. Technical Report 085, Smith College, July 2007. arXiv:0707.4258v2 [cs.CG].
- [IOV07b] Jin-ichi Itoh, Joseph O’Rourke, and Costin Vilcu. Unfolding convex polyhedra via quasigeodesics: Abstract. In *Proc. 17th Annu. Fall Workshop Comput. Comb. Geom.*, November 2007.
- [IOV09] Jin-ichi Itoh, Joseph O’Rourke, and Costin Vilcu. 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.
- [IV08a] Jin-ichi Itoh and Costin Vilcu. Criteria for farthest points on convex surfaces. *Mathematische Nachrichten*, to appear, 2008.
- [IV08b] Jin-ichi Itoh and Costin Vilcu. Geodesic characterizations of isosceles tetrahedra. Preprint, 2008.
- [Kob67] Shoschichi Kobayashi. On conjugate and cut loci. In S. S. Chern, editor, *Studies in Global Geometry and Analysis*, pages 96–122. Mathematical Association of America, 1967.
- [MP08] Ezra Miller and Igor Pak. Metric combinatorics of convex polyhedra: Cut loci and nonoverlapping unfoldings. *Discrete Comput. Geom.*, 39:339–388, 2008.
- [OV09] Joseph O’Rourke and Costin Vilcu. A new proof for star unfoldings of convex polyhedra. Manuscript in preparation, 2009.
- [Pog49] Aleksei V. Pogorelov. Quasi-geodesic lines on a convex surface. *Mat. Sb.*, 25(62):275–306, 1949. English transl., *Amer. Math. Soc. Transl.* 74, 1952.
- [Pog73] Aleksei V. Pogorelov. *Extrinsic Geometry of Convex Surfaces*, volume 35 of *Translations of Mathematical Monographs*. American Mathematical Society, Providence, RI, 1973.
- [Poi05] Henri Poincaré. Sur les lignes géodésiques des surfaces convexes. *Trans. Amer. Math. Soc.*, 6:237–274, 1905.
- [Sak96] Takashi Sakai. *Riemannian Geometry*. Translation of Mathematical Monographs 149. Amer. Math. Soc., 1996.

- [SS86] Micha Sharir and Amir Schorr. On shortest paths in polyhedral spaces. *SIAM J. Comput.*, 15:193–215, 1986.
- [SS08] Yevgeny Schreiber and Micha Sharir. An optimal-time algorithm for shortest paths on a convex polytope in three dimensions. *Discrete & Comput. Geom.*, 39:500–579, 2008.
- [ST96] Katsuhiro Shiohama and Minoru Tanaka. Cut loci and distance spheres on Alexandrov surfaces. *Séminaires & Congrès*, 1:531–559, 1996. Actes de la Table Ronde de Géométrie Différentielle (Luminy, 1992).