Vertex-Transplants on a Convex Polyhedron

Joseph O'Rourke

Abstract

Given any convex polyhedron $\mathcal{P}$ of sufficiently many vertices n, and with no vertex's curvature greater than π, it is possible to cut out a vertex, and paste the excised portion elsewhere along a vertex-to-vertex geodesic, creating a new convex polyhedron $\mathcal{P}'$. Although $\mathcal{P}'$ could have, in degenerate situations, as many as 2 fewer vertices, the generic situation is that $\mathcal{P}'$ has $n+2$ vertices. $\mathcal{P}'$ has the same surface area as $\mathcal{P}$, and the same total curvature but with some of that curvature redistributed.

1 Introduction

The goal of this paper is to prove the following theorem:

Theorem 1 *For any convex polyhedron $\mathcal{P}$ of $n > N$ vertices, none of which have curvature greater than π, there is a vertex v_0 that can be cut out along a digon of geodesics, and the excised surface glued to a geodesic on $\mathcal{P}$ connecting two vertices v_1, v_2. The result is a new convex polyhedron $\mathcal{P}'$ with, generically, $n+2$ vertices, although in various degenerate situations it could have $\{n-2, n-1, n, n+1\}$ vertices. $N = 16$ suffices.*

I conjecture that N can be reduced to 4 so that the theorem holds for all convex polyhedra with the stated curvature condition. Whether this curvature condition is necessary is unclear.

I have no particular application of this result, but it does raise several interesting questions (Sec. 8), including: Which convex polyhedra can be transformed into one another via a series of *vertex-transplants*?

2 Examples

Before detailing the proof, we provide several examples.

We rely on Alexandrov's celebrated gluing theorem [Ale05, p.100]: If one glues polygons together along their boundaries[1] to form a closed surface homeomorphic to a sphere, such that no point has more than 2π incident surface angle, then the result is a convex polyhedron, uniquely determined up to rigid motions. Although we use this theorem to guarantee that transplanting a vertex on $\mathcal{P}$ creates a new convex polyhedron $\mathcal{P}'$, there is as yet no effective procedure to actually construct $\mathcal{P}'$, except when $\mathcal{P}'$ has only a few vertices or special symmetries.

In the examples below, we use some notation that will not be fully explained until Sec. 3.

Cube. Fig. 1 shows excising a unit-cube corner v_0 with geodesics γ_1 and γ_2, each of length 1, and then suturing this digon into the edge v_1v_2. Although a paper model reveals a clear 10-vertex polyhedron (points x and y become vertices of $\mathcal{P}'$), I have not constructed it numerically.

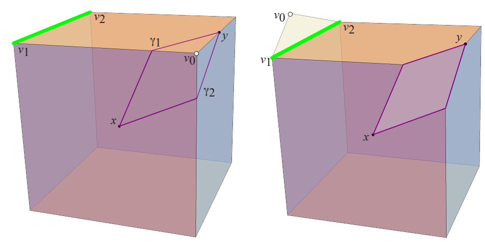

Figure 1: Left: Digon xy surrounding v_0. Right: v_0 transplanted to v_1v_2; v_0 is the apex of a doubly covered triangle, the digon flattened. Hole to be sutured closed to form $\mathcal{P}'$.

Regular Tetrahedron. Let the four vertices of a regular tetrahedron of unit edge length be v_0, v_1, v_2 forming the base, and apex v_3. Place a point x on the edge v_3v_0, close to v_3. Then one can form a digon starting from x and surrounding v_0 with geodesics γ_1 and γ_2 to a point y on the base, with $|\gamma_1| = |\gamma_2| = 1$. See Fig. 2. This digon can then be cut out and pasted into edge v_1v_2, forming an irregular 6-vertex polyhedron $\mathcal{P}'$.

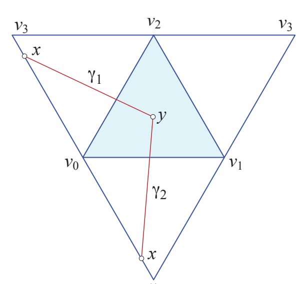

Figure 2: Unfolding of tetrahedron, apex v_3. Digon γ_i connect x to y, surrounding v_0.

[1] To "glue" means to identify boundary points.

Doubly Covered Square. Alexandrov's theorem holds for doubly covered, flat convex polygons, and vertex-transplanting does as well. A simple example is cutting off a corner of a doubly covered unit square with a unit length diagonal, and pasting the digon onto another edge. The result is another doubly covered polygon: see Fig. 3.

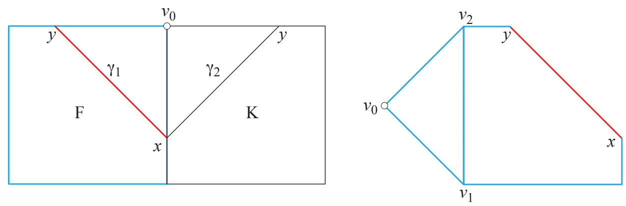

Figure 3: A doubly covered square $\mathcal{P}$ (front F, back K) converted to a doubly covered hexagon $\mathcal{P}'$.

A more interesting example is shown in Fig. 4. The indicated transplant produces a 6-vertex polyhedron $\mathcal{P}'$—combinatorially an octahedron—whose symmetries make exact reconstruction feasible. Vertices v_0 and v_3 retain their curvature π, and the remaining four vertices of $\mathcal{P}'$, v_1, v_2, x, y, each have curvature $\pi/2$.

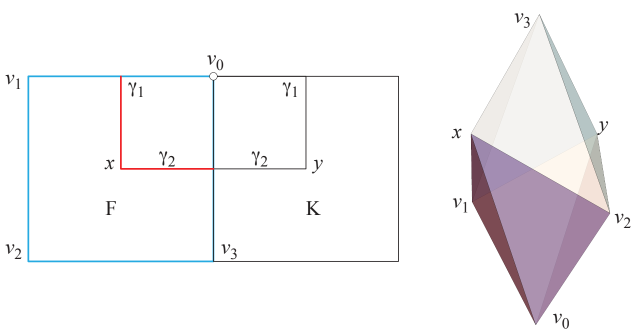

Figure 4: Transplanting v_0 to v_1v_2 on a doubly covered square (from F, back K) leads to a non-flat polyhedron $\mathcal{P}'$.

Doubly Covered Equilateral Triangle. The only polyhedron for which I am certain Theorem 1 (without restrictions) fails is the doubly covered, unit side-length equilateral triangle. The diameter $D = 1$ is realized by the endpoints of any of its three unit-length edges. Any other shortest geodesic is strictly less than 1 in length, as illustrated in Fig. 5. Thus there is no opportunity to create a digon of length 1 surrounding a vertex.

3 Preliminaries

Let the vertices of $\mathcal{P}$ be v_i, and let the curvature (angle gap) at v_i be ω_i. We assume all vertices are corners in the sense that $\omega_i > 0$. Let v_0, v_1, v_2 be three vertices,

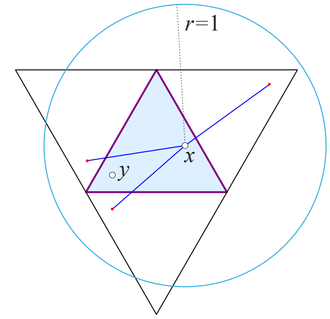

Figure 5: Point x is on the front, y on the back. Three images of y are shown, corresponding to the three paths from x to y. The shortest of these paths is never ≥ 1 unless both x and y are (different) corners.

labeled so that ω_0 is smallest, $\leq \omega_1, \omega_2$. Let v_1v_2 be the shortest geodesic on $\mathcal{P}$ connecting v_1 and v_2, with $|v_1v_2| = c$ its length. Often such a shortest geodesic is called a *segment*. We will show that, with careful choice of v_0, v_1, v_2, we can cut out a digon X of length c surrounding v_0, and paste it into v_1v_2 slit open. (A *digon* is a pair of shortest geodesics of the same length connecting two points on $\mathcal{P}$.)

The technique of gluing a triangle along a geodesic v_1v_2 on $\mathcal{P}$ was introduced by [Ale05, p.240], and employed in [OV14] to merge vertices. Excising a digon surrounding a vertex is used in [INV11, Lem.2]. What seems to be new is excising from one place on $\mathcal{P}$ and inserting elsewhere on $\mathcal{P}$.

Let $C(x)$ be the *cut locus* on $\mathcal{P}$ with respect to point $x \in \mathcal{P}$. (In some computer science literature, this is called the *ridge tree* [AAOS97].) $C(x)$ is the set of points on $\mathcal{P}$ with at least two shortest paths from x. It is a tree composed of shortest paths; in general, each vertex of $\mathcal{P}$ is a degree-1 leaf of the closure of $C(x)$.

We will need to exclude positions of x that are non-generic in that $C(x)$ includes one or more vertices. It was shown in [AAOS97, Lem.] that the surface of $\mathcal{P}$ may be partitioned into $O(n^4)$ *ridge-free* regions, determined by overlaying the cut loci of all vertices: $\bigcup_i C(v_i)$. Say that $x \in \mathcal{P}$ is *generic* if it lies strictly inside a ridge-free region. For later reference, we state this lemma:

Lemma 2 *Within every neighborhood of any point $x \in \mathcal{P}$, there is a generic $y \in P$.*

Proof. This follows because ridge-free regions are bounded by cut-loci arcs, each of which is a 1-dimensional geodesic. □

For generic x, the cut locus in the neighborhood of a vertex v_0 consists of a geodesic segment s open at v_0 and continuing for some positive distance before reaching a junction u of degree-3 or higher. Let $\delta(x, u) = \delta$ be the length of s; see Fig. 6.

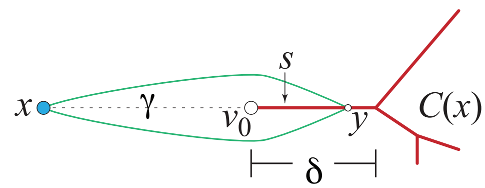

Figure 6: Geodesic segment s of $C(x)$ (red) incident to vertex v_0. A pair of shortest geodesics from x to s are shown (green).

4 Surgery Procedure

We start with and will describe the procedure for any three vertices v_0, v_1, v_2, but later (Sec. 5) we will chose specific vertices.

Let x be a generic point on $\mathcal{P}$ and γ a shortest geodesic to v_0 with length $|\gamma| = |v_1v_2| = c$. The existence of such an x is deferred to Sec. 5. If we move x along γ toward v_0, γ splits into two geodesics γ_1, γ_2 connecting x to a point $y \in C(x)$. If we move x a small enough distance ε, then y will lie on the segment $s \subset C(x)$ as in Fig. 6. Because Lemma 2 allows us to choose x to lie in a ridge-free region R, we can ensure that s has a length $|s| = \delta > 0$. Now γ_1, γ_2 form a digon X surrounding v_0. With sufficiently small ε, we can ensure that X is empty of other vertices. During the move of x along γ, we can at all times maintain that $|\gamma_1| = |\gamma_2| = c$, as illustrated in Fig. 7.

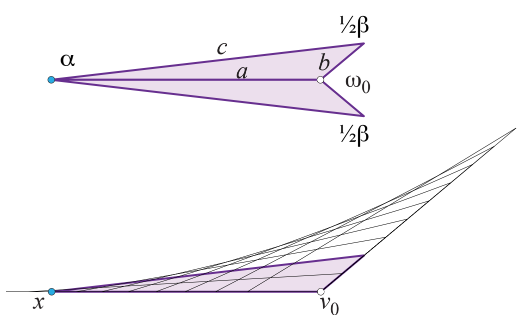

Figure 7: Sliding x along γ toward v_0 while maintaining length c constant.

Let the digon angles at x and at y be α and β respectively. By Gauss-Bonnet, we have $\alpha + \beta = \omega_0$:

$$\tau + \omega_0 = 2\pi = ((\pi - \alpha) + (\pi - \beta)) + \omega_0 = 2\pi \ ,$$

where the turn angle τ is only non-zero at the endpoints x and y. In particular, $0 < \alpha, \beta < \omega_0$. These inequalities are strict because the digon wraps around v_0 after moving x toward v_0, so $\alpha > 0$.

Now we can suture-in the digon X to a slit along v_1v_2 because:

- The lengths match: $|v_1v_2| = c$ and $|\gamma_1| = |\gamma_2| = c$.
- The curvatures at v_1, v_2 remain positive: $\alpha, \beta < \omega_0 \le \omega_1, \omega_2$, so $\omega_1 - \alpha > 0$ and $\omega_2 - \beta > 0$.

We then close up the digon on the surface of $\mathcal{P}$ and invoke Alexandrov's theorem to obtain $\mathcal{P}'$. We now detail the curvature consequences at the five points involved in the surgery: v_0, v_1, v_2, x, y.

- v_0 is unaltered, just moved, i.e, transplanted.
- If x and/or y were not vertices before the transplant, they become vertices after the transplant, of curvatures α and β respectively.
- If x and/or y were vertices, they remain vertices with larger curvatures.
- Because $\alpha < \omega_0 \le \omega_1$, the change at v_1 cannot flatten v_1. So v_1 remains a vertex, as does v_2.

So the new polyhedron $\mathcal{P}'$ has n, $n+1$, or $n+2$ vertices, depending on whether x and/or y was already a vertex.

We note that the condition that $\omega_0 \le \omega_1, \omega_2$ is in fact more stringent that what is required to ensure that the curvatures at v_1, v_2 remain non-negative. The latter implies that $\omega_0 \le \omega_1 + \omega_2$, a considerably weaker condition. And indeed, reversing the transplant may not satisfy $\omega_0 \le \omega_1, \omega_2$. For example, in the cube-corner transplant in Fig. 1, $\alpha = 41°$, $\beta = 49°$, $\omega_0 = 90°$. So reversing, x, y play the roles of v_1, v_2, with curvatures $41°, 49°$, both less than ω_0. This shows that $\omega_0 \le \omega_1, \omega_2$ is not necessary; but with that condition, we can guarantee a transplant.

5 Existence of v_0, v_1, v_2

In order to apply the procedure just detailed, we need several conditions to be simultaneously satisfied:

(1) $\omega_0 \le \omega_1, \omega_2$.

(2) $|v_1v_2| = |\gamma_1| = |\gamma_2| = c$.

(3) v_1v_2 should not cross the digon X.

Although (1) is satisfied by any three vertices, just by identifying v_0 with the smallest curvature, the difficulty is that if v_1v_2 is long—say, realizing the diameter of $\mathcal{P}$—then we need there to be an equally long geodesic from x to v_0, to satisfy (2). A solution is to choose v_1 and v_2 to be the nearest neighbors on $\mathcal{P}$, so that $|v_1v_2|$ is small. But then if ω_1, ω_2 are both small, we may not be able to locate a v_0 with a smaller ω_0. We resolve these tensions as follows:

1. We choose v_0 to be a vertex with minimum curvature, so automatically $\omega_0 \le \omega_1, \omega_2$ for any choices for v_1 and v_2.
2. Several steps to achieve (2):

(a) We choose v_1, v_2 to achieve the smallest nearest-neighbor distance $\mathrm{NN}_{\min} = r$ over all pairs of vertices (excluding v_0), so v_1v_2 is as short as possible.
(b) We prove that the nearest neighbor distance r satisfies $r < \frac{1}{2}D$, where D is the diameter of $\mathcal{P}$.
(c) We prove that there is an x such that $d(x, v_0) \geq \frac{1}{2}D$.

Together these imply that we can achieve $|v_1v_2| = |\gamma_1| = |\gamma_2|$.

3. We show that if v_1v_2 crosses X, then another point point x may be found that avoids the crossing. This last point is the only use of the assumption that $\omega_i \leq \pi$ for all vertices v_i.

The next section addresses items (1) and (2) above, and Sec. 7 addresses item (3).

6 Relationship to Diameter D

The *diameter* $D(\mathcal{P})$ of $\mathcal{P}$ is the length of the longest shortest path between any two points. The lemma below ensures we can find a long-enough geodesic $\gamma = xv_0$.

Lemma 3 *For any $x \in \mathcal{P}$, the distance ρ to a point $f(x)$ furthest from x is at least $\frac{1}{2}D$, where $D = D(\mathcal{P})$ is the diameter of $\mathcal{P}$.*

Proof. [2] Let points $y, z \in \mathcal{P}$ realize the diameter: $d(y, z) = D$. For any $x \in \mathcal{P}$,

$$D = d(y, z) \leq d(y, x) + d(x, z)$$

by the triangle inequality on surfaces [Ale06, p.1]. Also we have $\rho \leq d(x, y)$ and $\rho \leq d(x, z)$ because ρ is the furthest distance. So $D = d(y, z) \leq 2\rho$, which establishes the claim. □

Next we establish that the smallest distance (via a shortest geodesic) between a pair of vertices of $\mathcal{P}$, $\mathrm{NN}_{\min}$—the *nearest neighbor distance*—cannot be large with respect to the diameter $D = D(\mathcal{P})$.

6.1 Nearest-Neighbor Distance

Here our goal is to show that sufficiently many points on P cannot all have large nearest-neighbor (NN) distances. First we provide two examples.

1. Let P be a regular tetrahedron with unit edge lengths. D is determined by a point in the center of the base connecting to the apex, so $D = \frac{4}{3}\frac{\sqrt{3}}{2} = \frac{2}{\sqrt{3}}$. The NN-distance is $1 = \frac{\sqrt{3}}{2}D = 0.866\,D$.

[2] Proof suggested by Alexandre Eremenko. `https://mathoverflow.net/a/340056/6094`. See also [IRV19].

2. Let P be a doubly covered regular hexagon, with unit edge lengths. Then $D = 2$, connecting opposite vertices, and the NN-distance is $1 = \frac{1}{2}D$.

Our goal is to ensure the NN distance is at most $\frac{1}{2}D$, which is not achieved by the regular tetrahedron but is for the hexagon. We achieve this by insisting $\mathcal{P}$ have many vertices.

Lemma 4 *Let $\mathcal{P}$ be a polyhedron with diameter D. Let S be a set of distinguished points on $\mathcal{P}$, with $|S| \geq N$. Let r be the smallest* NN-distance *between any pair of points of S. Then $r < D/(\sqrt{N}/2)$. In particular, for $N = 16$, $r < \frac{1}{2}D$.*

Proof.

1. Let a geodesic from x to y realize the diameter D of $\mathcal{P}$. Let U be the *source unfolding* of $\mathcal{P}$ from source point x [DO07, Chap.24.1.1]. U does not self-overlap, and fits inside a circle of radius D; see Fig. 8. Thus the surface area of $\mathcal{P}$ is at most πD^2.

2. Let r be the smallest NN-distance, the smallest separation between a pair of points in S. Then disks of radius $r/2$ centered on points of S have disjoint interiors. For suppose instead two disks overlapped. Then they would be separated by less than r, a contradiction.

3. N non-overlapping disks of radius $r/2$ cover an area of $N\pi(r/2)^2$, which must be less than[3] the surface area of $\mathcal{P}$:

$$N\pi(r/2)^2 < \pi D^2 \quad (1)$$

$$r < \frac{D}{\sqrt{N}/2} \quad (2)$$

Thus, for $N = 16$, $r < \frac{1}{2}D$.

□

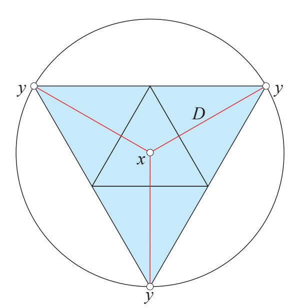

Figure 8: Source unfolding of a regular tetrahedron. xy realizes D.

[3] Strictly less than because disk packings leave uncovered gaps.

7 Crossing Avoidance

Although Lemma 3 ensures that we can find an $x = f(v_0)$ far enough from v_0 so that we can match $|\gamma|$ with $|v_1v_2|$, if γ crosses v_1v_2, the procedure in Sec. 4 fails. We now detail a method to locate another x in this circumstance. We partition crossings into several cases.

Recall that v_0 was excluded from the NN calculation of r, so v_0 could be closer to v_1 and/or v_2 than $r = |v_1v_2|$.

Case (1). Case: $d(v_0, v_i) > r$ for either $i = 1$ or $i = 2$. Assume $d(v_0, v_2) > r$. Then choose $\gamma = v_0v_2$. We can locate x near v_2 on γ to achieve $|xv_0| = r$. See Fig. 9.

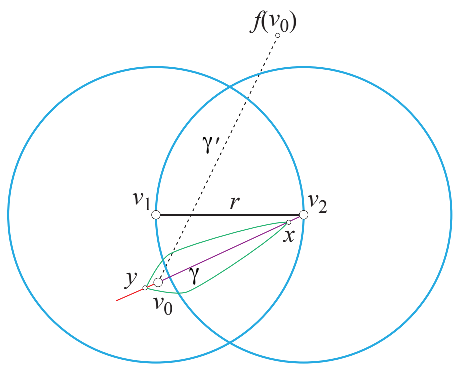

Figure 9: Crossing avoided: $d(v_0, v_2) > r$ (v_0 is outside v_2's r-disk).

Case (2). If $d(v_0, v_i) \leq r$ for $i = 1, 2$, then v_0 is located in the half-lune to the opposite side of (below) v_1v_2 from $f(v_0)$. It is possible that with large curvatures ω_1 and ω_2 that there is no evident "room" below v_1v_2 to locate an x far enough away so that $d(x, v_0) \geq r$. However, with assumptions on the maximum curvature per vertex, room can be found.

First we assume $\omega_i \leq \pi/2$ for all i. As Fig. 10 illustrates, it is possible to find a horizontal (parallel to v_1v_2) segment xv_0 either left or right of v_0. In the figure, $\omega_1 = \omega_2 = \pi/2$, with segment $(v_0, f(v_0))$ slanting to the right, which requires the angle gap to interfere with connecting to v_0 to the right. But then to the left there is room for an x with $|xv_0| = r$.

Case (3). For larger curvatures, there might not be room either right or left for an x achieving $|xv_0| = r$.[4] Indeed, in the most extreme case, the situation could resemble a doubly covered equilateral triangle with $\omega_i = \frac{4}{3}\pi$, which we saw in Sec. 2 violates Theorem 1. However, if we assume $\omega_i \leq \pi$ for all i, a long-enough γ to v_0 can be found.

[4] I have not found an argument that finds such an x below v_0.

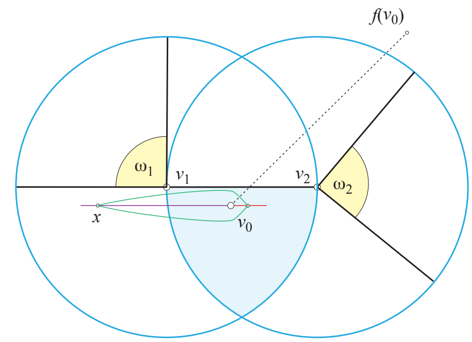

Figure 10: Curvatures $\leq \pi/2$ allows room for xv_0.

Assume the worst case, $\omega_1 = \omega_2 = \pi$. As illustrated in Fig. 11, there is neither "room" right nor left for a segment of length r incident to v_0, and with enough twisting at v_1 and v_2, no room below either. However, an r-long segment left of v_0 re-enters above v_1v_2 (red), and similarly right of v_0 (green). In fact, it is easy to see that the red and green segments above and below have total length $2r$, regardless of the orientation of the semicircle bounding the angle-gap lines through v_1 and v_2. So there is always enough room to locate x above v_1v_2 connecting "horizontally" to v_0 below.

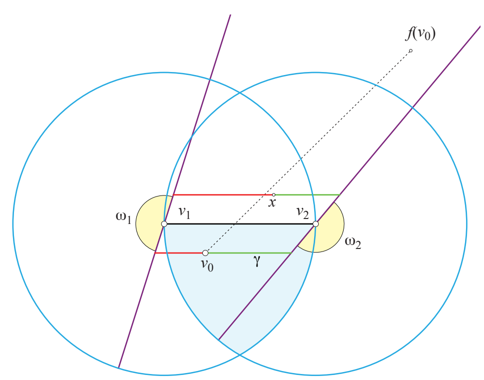

Figure 11: Crossing avoided: Both the red and green segments have total length r each.

8 Open Problems

1. Extend Theorem 1 to all convex polyhedra, i.e., lower $N = 16$ to $N = 4$, and remove the $\omega_i \leq \pi$ restriction.

2. Establish conditions that allow more freedom in the selection of the three vertices v_0, v_1, v_2. Right now, Thm. 1 requires following the restrictions detailed

in Sec. 5, but as we observed, these restrictions are not necessary for a successful transplant.

3. Study doubly covered convex polygons as a special case. When does a vertex transplant on a doubly covered polygon produce another doubly covered polygon? See again Sec. 2. (There is a procedure for identifying flat polyhedra [O'R10]; and see [INV11, Lem.4].)

4. What happens under repeated vertex-transplanting? Because the number of vertices n does not diminish except in degenerate situations, the procedure usually can be repeated indefinitely. Note that because $\alpha, \beta < \omega_0$, new smaller-curvature vertices are created at x and y.

5. Which convex polyhedra can be connected by a series of vertex-transplants? Recall that each vertex transplant is reversible, so it seems possible to connect two polyhedra $\mathcal{P}_1, \mathcal{P}_2$ via some canonical form $\mathcal{P}_c$, reversing the $\mathcal{P}_2$ transplants: $\mathcal{P}_1 \to \mathcal{P}_c \to \mathcal{P}_2$.

6. Can Thm. 1 be generalized to transplant several vertices within the same digon? For example, one can excise both endpoints of an edge of a unit cube with a digon of length $\sqrt{2}$ and suture that into a face diagonal.

7. Does the transplant guaranteed by Thm. 1 always increase the volume of $\mathcal{P}$? Note that a transplant flattens v_1 and v_2 by α and β, and creates new smallest curvature vertices, $\alpha, \beta < \omega_0$. So the overall effect seems to "round" $\mathcal{P}$.

Acknowledgement. I benefitted from the advice of Costin Vîlcu.

References

[AAOS97] Pankaj K. Agarwal, Boris Aronov, Joseph O'Rourke, and Catherine A. Schevon. Star unfolding of a polytope with applications. *SIAM J. Comput.*, 26:1689–1713, 1997.

[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.

[Ale06] Aleksandr D. Alexandrov. Intrinsic geometry of convex surfaces. In S. S. Kutateladze, editor, *A. D. Alexandrov: Selected Works: Part II*, pages 1–426. Chapman & Hall, Boca Raton, 2006.

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

[INV11] Jin-ichi Itoh, Chie Nara, and Costin Vîlcu. Continuous flattening of convex polyhedra. In *Centre de Recerca Matemàtica Documents*, volume 8, pages 95–98, 2011. XIV Spanish Meeting on Computational Geometry.

[IRV19] Jin-ichi Itoh, Jöel Rouyer, and Costin Vîlcu. Some inequalities for tetrahedra. `https://arxiv.org/abs/1906.11965`, 2019.

[O'R10] Joseph O'Rourke. On flat polyhedra deriving from Alexandrov's theorem. `http://arxiv.org/abs/1007.2016v2`, July 2010.

[OV14] Joseph O'Rourke and Costin Vîlcu. Development of curves on polyhedra via conical existence. *Comput. Geom.: Theory & Appl.*, 47:149–163, 2014.