Finding Weakly Simple Closed Quasigeodesics on Polyhedral Spheres

Jean Chartier ✉
Univ. Paris Est Creteil, CNRS, LAMA, F-94010 Creteil, France

Arnaud de Mesmay ✉
LIGM, CNRS, Univ. Gustave Eiffel, ESIEE Paris, F-77454 Marne-la-Vallée, France

──── **Abstract** ────

A closed quasigeodesic on a convex polyhedron is a closed curve that is locally straight outside of the vertices, where it forms an angle at most π on both sides. While the existence of a simple closed quasigeodesic on a convex polyhedron has been proved by Pogorelov in 1949, finding a polynomial-time algorithm to compute such a simple closed quasigeodesic has been repeatedly posed as an open problem. Our first contribution is to propose an extended definition of quasigeodesics in the intrinsic setting of (not necessarily convex) polyhedral spheres, and to prove the existence of a weakly simple closed quasigeodesic in such a setting. Our proof does not proceed via an approximation by smooth surfaces, but relies on an adapation of the disk flow of Hass and Scott to the context of polyhedral surfaces. Our second result is to leverage this existence theorem to provide a finite algorithm to compute a weakly simple closed quasigeodesic on a polyhedral sphere. On a convex polyhedron, our algorithm computes a simple closed quasigeodesic, solving an open problem of Demaine, Hersterberg and Ku.

1 Introduction

A geodesic is a curve on a surface, or more generally in a manifold, which is locally shortest. The study of geodesics on surfaces dates back at least to Poincaré [20] and led to a celebrated theorem of Lyusternik and Schnirelmann [18] proving that any Riemannian sphere admits at least three distinct simple (i.e., not self-intersecting) closed geodesics (while the initial proof of the theorem was criticized, the result is now well-established, see for example Grayson [15]). This bound is tight, as showcased by ellipsoids.

In this article, we investigate closed geodesics in a polyhedral setting. In such a setting, the following relaxed notion is key: a *quasigeodesic* is a curve such that the angle is at most π on both sides at each point of the curve. In 1949, Pogorelov [19] proved the existence of three simple (i.e., non self-intersecting) and closed quasigeodesics on any convex polyhedron. The proof is non-constructive and it was asked by Demaine and O'Rourke [13, Open Problem 24.24] whether one could compute such a closed quasigeodesic in polynomial time. Recent progress on this question was made by Demaine, Hersterberg and Ku [12] who provided the first algorithm to compute a closed quasigeodesic on a convex polyhedron, and their
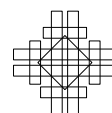
algorithm runs in pseudo-polynomial time. However, their algorithm is ill-adapted to find closed quasigeodesics which are simple – this has remained an open problem [12, Open Problem 1]. Furthermore, as they note, for this problem, "even a finite algorithm is not known or obvious": indeed there is no known upper bound on the combinatorial complexity of a simple closed quasigeodesic (for example the number of times that it intersects each edge), so there is no natural brute-force algorithm. We refer to the extensive introduction of [12] for a panorama on the difficulties in finding closed quasigeodesics.

Our results. Our contributions in this article are two-fold.

First, we extend the theorem of Pogorelov to a non-convex and non-embedded setting. Precisely, we work in the abstract setting of compact *polyhedral spheres*, which consist of the following data: (1) a finite collection of Euclidean polygons, and (2) gluing rules between pairs of boundaries of equal length, so that the topological space resulting from the gluings is a topological sphere. A face, edge or vertex of a polyhedral sphere is respectively a polygon, an edge or a vertex of one of the polygons, and a vertex is *convex* (respectively *concave*) if the sum of the angles of the polygons around the vertex is at most 2π, respectively at least 2π. Let us emphasize that such a polyhedral sphere is not a priori embedded in $\mathbb{R}^3$. In particular, edges of the triangles might not be shortest paths. This intrinsic description of non-smooth surfaces appears under various names in the literature, see, e.g., piecewise-linear surfaces [14] or intrinsic triangulations [21], and dates back to at least Alexandrov, who proved [2, Chapter 4] that when all the vertices are convex, such a polyhedral sphere is the metric structure of a unique convex polyhedron in $\mathbb{R}^3$ (see [17] for an algorithmic version of this result). In the non-convex case, a celebrated theorem of Burago and Zalgaller [8], shows that one can always find a piecewise-linear isometric embedding of a compact polyhedral sphere into $\mathbb{R}^3$, but it might require a large number of subdivisions and the proof has to our knowledge not been made algorithmic.

Note that by definition, a polyhedral sphere is locally Euclidean at every point that is not a vertex. We propose the following generalization of the definition of quasigeodesics to a polyhedral sphere S: a closed quasigeodesic is a closed curve that is locally a straight line around any point that is not a vertex, and that is locally a pair of straight lines around a vertex, forming an angle *at most* π on each side if the vertex is convex, and forming an angle *at least* π on each side if the vertex is concave. A closed curve $\gamma : \mathbb{S}^1 \to S$ is *simple* if it is injective, and is *weakly simple* if it is a limit of simple curves (see Section 2 for details).

Our first theorem shows the existence of a weakly simple closed quasigeodesic of controlled length on a polyhedral sphere. We denote by M the *edge-sum* of S, which we define as the sum of the lengths of the edges of an iterated barycentric subdivision of a triangulation of S.

▶ **Theorem 1** (Existence). *Let S be a polyhedral sphere and denote by M its edge-sum. There exists a weakly simple closed quasigeodesic of length at most M.*

The original proof of Pogorelov in the convex case proceeds by first approximating the polyhedron with smooth surfaces, and then taking the limit of the simple closed geodesics on the smooth surfaces, whose existence is guaranteed by the Lyusternik–Schnirelmann theorem. The proof technique for that latter theorem, originating from the work of Birkhoff [5], goes roughly as follows: we consider *sweep-outs*, i.e., a family of simple closed curves sweeping the polyhedron from one point to another point (see Section 2 for a precise definition), and consider the sweep-out where the longest curve has minimal length. Then, by applying a *curve-shortening process*, one can use this optimal sweep-out to find simple closed geodesics. This last step is notoriously perilous [3, 4, 15], hence the tumultuous history of the Lyusternik-Schnirelmann theorem. Our proof proceeds by working directly on the polyhedral sphere

and we prove the existence of a weakly simple closed quasigeodesic using a similar technique based on sweep-outs. Our key technical contribution is to rely on a curve-shortening process that is well-adapted to the polyhedral structure of the problem: we adapt the *disk flow* originally designed by Hass and Scott [16] for Riemannian surfaces so as to handle the disks formed by the stars of vertices in a seamless way. We are hopeful that this polyhedral variant of the disk flow could find further applications in the study of quasigeodesics.

Theorem 1 provides, in addition to the existence of a weakly simple closed quasigeodesic, a bound on its length. Our second result is to leverage this bound in order to control the combinatorics of the quasigeodesic, which allows us to design a finite algorithm to compute a weakly simple closed quasigeodesic on a polyhedral sphere.

▶ **Theorem 2** (Algorithm). *Given a polyhedral sphere S, we can compute a weakly simple closed quasigeodesic in time exponential in n and $\lceil M/h \rceil$, where n is the number of vertices of S, M is its edge-sum, and h is the smallest altitude over all triangles of some triangulation of S.*

Note that a bound on the length of a quasigeodesic does not translate directly into a bound on the number of times that it crosses each edge of the polyhedral sphere, as these crossings could happen arbitrarily close to vertices, and thus contribute an arbitrarily small length. Our proof of Theorem 2 investigates the local geometry of quasigeodesics around vertices to show that this does not happen too much, and that one can indeed bound the multiplicity of each edge. Then, our algorithm guesses the correct combinatorics of the simple closed quasigeodesic and checks in polynomial time that it is realizable.

Our proof techniques for Theorem 1 only provide the existence of weakly simple quasi-geodesics instead of simple quasigeodesics. We believe this to be a necessary evil in any generalization to the non-convex case, as shortest paths accumulate on concave vertices, making it impossible to define a curve-shortening process in the neighborhood of those which preserves simplicity. However, when all the vertices are convex, the result of Pogorelov does show the existence of a (actually three) simple closed quasigeodesics, where we include as a degenerate simple case a curve connecting twice two vertices of curvature at least π. Furthermore, his proof also provides an upper bound on the length of this simple quasigeodesic, as we explain an the end of Section 4. Since our algorithm behind Theorem 2 only relies on such an upper bound on the length and on the (weak) simplicity of the sought after curve, we can also use it to compute simple closed quasigeodesics in the convex case. This solves Open Problem 1 of [12], but note that we are still a long way off a polynomial-time algorithm.

Some of the proofs have been omitted and are available in the full version [11].

2 Preliminaries

In this article, a *polyhedral sphere* is a finite collection of Euclidean polygons, and gluing rules for boundaries of the same length, so that the space obtained by identifying the boundaries of the polygons via the gluing rules is homeomorphic to a sphere. Such a sphere is naturally endowed with a metric which is locally Euclidean at every point except at the vertices of the polygons, where it might display a *conical singularity*: if the total angle of the polygons glued around that vertex is larger than 2π (respectively at most 2π), we say that the vertex is *concave* (respectively *convex*), and its *curvature* is the angular defect compared to 2π (which is thus negative for concave vertices). Given a (not necessarily convex) polyhedron described via the coordinates of its vertices in $\mathbb{R}^3$, one can easily compute the underlying polygons and thus the structure as a Euclidean sphere. The reverse direction of embedding a

polyhedral sphere in $\mathbb{R}^3$ is significantly more intricate (see [17] for the convex case and [8] for the general case), hence our choice of the intrinsic model.

Triangulating each polygon defining a polyhedral sphere yields a *triangulated polyhedral sphere*. Furthermore, by doing up to two barycentric subdivisions in each triangle if necessary, we can assume that there are no loops nor multiple edges in this triangulation. Note that this triangulation and these subdivisions do not change the metric of the sphere and do not impact quasigeodesicity (see next paragraph). Therefore, for convenience, in this article we will always assume that our polyhedral spheres are triangulated and that they contain neither loops nor multiple edges, and we will denote such a sphere by S from now on. A *shelling* of a triangulated sphere S is an order $(T_1, \ldots, T_\ell)$ on the triangles that S consists of so that for all $i \in [1, \ell-1]$, $\bigcup_{k=1}^{k=i} T_k$ is homeomorphic to a 2-disk D^2. It is well-known that all the triangulated spheres are shellable, for example because, by Steinitz's theorem [22, Chapter 4] they form the 2-skeleton of a polytope, and those are shellable [6]. Throughout this article, we use the following notations for a polyhedral sphere: its vertices are denoted by $p_1, \ldots, p_n$, its edges by $e_1, \ldots, e_m$ (or sometimes e_{ij} to emphasize the vertices that it connects to) and its triangles by $T_1, \ldots, T_\ell$. The order induced by the numbering of the triangles is a shelling order. The *star* of vertex p_i, denoted by $\mathcal{C}_i$ is the union of the triangles T_k having p_i for common vertex, identified along the edges adjacent to p_i. It is *convex* (resp. *concave*) if p_i is (but note that the shortest path in S between two points of a convex star is not necessarily contained in that star). We optionally rename the vertices of P to have $p_1 \in T_1$ and $p_n \in T_\ell$. Finally, we denote by M the sum of the lengths of the edges of S, and by h the smallest altitude of all the triangles in S. Note that h is a lower bound on the distance between any two vertices. For γ an edge or a curve on S, we denote by $L(\gamma)$ its length.

A *closed curve c* on S is a continuous map $c : \mathbb{S}^1 \to S$. A closed curve is *piecewise-linear* if it is locally straight except at a finite number of points.

▶ **Definition 3.** *A closed curve is a* quasigeodesic *if it is locally straight around every point of S that is not a vertex, and around a vertex it forms an angle at most (respectively at least) π on both sides if the vertex is* convex *(respectively concave).*

We emphasize that this definition is non-standard in the non-convex case, where it is sometimes simply forbidden for a quasigeodesic to go through a concave vertex [13]. Note that a quasigeodesic is straight around a vertex with zero curvature. A closed curve is *simple* if it is injective. Throughout this article, all the curves will always be parameterized at constant speed. We endow the space of piecewise-linear curves with the uniform convergence metric, i.e., $d(c_1, c_2) = \max_{t \in \mathbb{S}^1} d(c_1(t), c_2(t))$. A closed curve is *weakly simple* if it is a limit of simple curves for this metric: intuitively a weakly simple curve is a curve with tangencies but no self-crossings. We denote by $\mathcal{P}$ the set of constant closed curves, i.e., closed curves c such that there exists $p \in S$ such that $\forall t \in \mathbb{S}^1, c(t) = p$.

We denote by Ω the space of rectifiable closed curves of length at most M. This space is compact for the uniform convergence metric, as can be shown via the Arzelà-Ascoli theorem, the bound on the length and the constant-speed parameterization providing equicontinuity (see for example [7, Theorem 2.5.14]). We denote by Ω^{pl} the subspace of Ω consisting of piecewise-linear and weakly simple closed curves. A *monotone sweep-out* of S is a continuous map $\beta : \mathbb{S}^2 \longrightarrow S$, where $\mathbb{S}^2$ is seen as the quotient of the cylinder $[0, 1] \times \mathbb{S}^1$ by the relation which identifies the circles $(0, \mathbb{S}^1)$ and $(1, \mathbb{S}^1)$ to two points, and such that:

- $\beta(0, \cdot)$ and $\beta(1, \cdot)$ belong to $\mathcal{P}$, i.e., are two constant closed curves on S,
- β has topological degree one,

- for $s \in (0,1)$, each *fiber* $\beta(s, \cdot) : \mathbb{S}^1 \longrightarrow S$ belongs to Ω^{pl}, and
- the sweep-out is monotone, i.e., if D_s denotes the disk to the left of $\beta(s, \cdot)$, the disks D_s are nested: $D_s \subseteq D_{s'}$ for $s' > s$.

The requirement on the topological degree informally means that each point is covered once by the sweep-out ; it is there to prevent trivial sweep-outs (for example constant at a point). It can be replaced by the requirement that the starting and endpoints are distinct. The monotonicity corresponds to the third condition, and typical sweep-outs in the literature do not assume it (see [9]), but in this paper we will only use monotone sweep-outs and thus for simplicity we will henceforth drop the word monotone. The *width* of a sweep-out is the length of the longest fiber. We denote by $\mathcal{B}$ the space of sweep-outs.

The algorithm underlying Theorem 2 has complexity exponential in $\lceil M/h \rceil$, i.e., it depends on the actual values of the lengths of the boundaries of the polygons. Therefore, we do not work on a real RAM model and rely rather on a word RAM model, which is powerful enough to express all the operations that we require: see for example [12, Section 2] for a description of the $O(1)$-Expression RAM model which can be encoded in the word RAM model and allows for a restricted notion of real numbers and algebraic operations thereon.

3 Disk flow and sweep-outs

We start by describing a monotone sweep-out of controlled width.

▶ **Lemma 4.** *Let S be a triangulated polyhedral sphere of edge-sum M. There exists a monotone sweep-out of S of width at most M.*

This lemma is proved by sweeping triangles one by one, in the order prescribed by a shelling order of the sphere S.

The disk flow. We define here a curve-shortening process that we call the disk flow, which is an iterative process Φ shortening locally a curve in Ω^{pl} successively in each star $\mathcal{C}_i$, with the key property that the only fixed points of Φ are quasigeodesics or trivial curves. In a second step, we will extend Φ into a map $\hat{\Phi}$ that acts on monotone sweep-outs, which will require interpolating at the points where Φ is discontinuous. This disk flow is directly inspired by the work of Hass and Scott [16] who defined an analogous flow on Riemannian surfaces. The key difference with their setup is that the star $\mathcal{C}_i$ around a convex vertex is not strongly convex (i.e. there is no uniqueness of shortest paths), which causes additional tears when extending Φ to sweep-outs and thus requires further operations. Furthermore, instead of working with very small convex disks as they are doing, we work directly with the stars $\mathcal{C}_i$ as we strive to preserve curves whose piecewise-linear structure matches that of S. This requires us to deal with tangencies with the boundaries of stars in a different manner.

Let c be a curve in Ω^{pl} and let $\mathcal{C}_i$ be a star crossed by $c \in \Omega^{pl}$. An *arc* of $\mathcal{C}_i$ is a restriction of c whose image is a connected component of $\mathcal{C}_i \cap Im(c)$. Let γ be an arc of $\mathcal{C}_i$, from a closed curve c. The points $\gamma(t_0) = c(t_0) \in \partial\mathcal{C}_i$ such that $c([t_0 - \varepsilon, t_0))$ or $c((t_0, t_0 + \varepsilon])$ is contained in the interior of $\mathcal{C}_i$ for a small enough $\varepsilon > 0$, are called the *gates* of γ. Note that two kinds of arcs have no gates: loops strictly inside the star and arcs never meeting the interior of the star. Unless γ is included in $\mathcal{C}_i$, the orientation of $\mathbb{S}^1$ naturally designates a first gate, denoted by **front**(γ), and a final gate, denoted by **exit**(γ). The gates can give access to the interior of the star for values of t greater (resp. less) than t_0 – we say that the gate is open to the right (resp. to the left). A gate can be open to the right and to the left. Thus, front gates are open to the right and exit gates are open to the left. Figure 1 illustrates different possible sequences of gates.
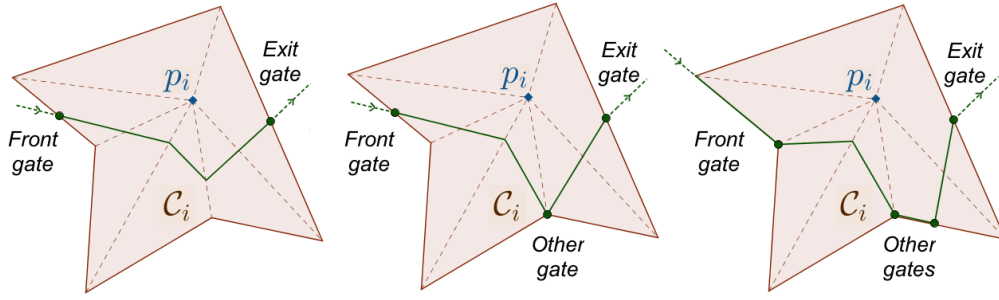
Figure 1 Three examples of sequences of gates.

Relative to two gates A and B *and independently of the path followed between A and B*, we define the **right region** $\mathcal{C}_i^r(A, B)$ and the **left region** $\mathcal{C}_i^\ell(A, B)$ of the star, as being the two parts of $\mathcal{C}_i$ whose union is $\mathcal{C}_i$ and which intersect along the edges $[Ap_i]$ and $[p_iB]$. The orientation right/left is chosen compatible with that of c between the two gates. The angles of the regions at the p_i vertex are called the **right angle** $\theta_r(A, B)$ and the **left angle** $\theta_\ell(A, B)$.

▶ **Lemma 5.** *Let c be a curve in Ω^{pl}. There exists a map $\Phi : \Omega^{pl} \longrightarrow \Omega^{pl}$ satisfying the following two conditions :*
- *The only fixed points of Φ are quasigeodesics and constant curves.*
- *$L(\Phi(c)) \leq L(c)$, with equality if and only if c is a fixed point.*

We stress that the map Φ is in general *not* continuous.

Proof. We define Φ as follows. Let c be a closed curve in Ω^{pl}. We pick an arbitrary order on the vertices of S, which induces an arbitrary order on the stars $\mathcal{C}_i$. The map Φ consists in repeating in this order a straightening process Φ_{loc}^i successively in each star. Consider in $\mathcal{C}_i$ an arc γ of c. Note that between two of its consecutive gates, A open to the right and B open to the left, γ lies in $\mathcal{C}_i$.

If $\mathcal{C}_i$ is convex, the straightening is defined as follows for each subset of γ between two consecutive gates (which by a slight abuse of notation we also denote by γ):
- If $p_i \in \gamma$ and if $\theta_r(A, B)$ and $\theta_\ell(A, B)$ are less than or equal to π, we replace γ by $[Ap_i] \cup [p_iB]$.
- If $p_i \notin \gamma$ and if $\theta_r(A, B)$ and $\theta_\ell(A, B)$ are less than or equal to π, we replace γ by the shortest path between A and B staying in the same region relative to A and B.
- If $\theta_r(A, B)$ (resp. $\theta_\ell(A, B)$) is strictly greater than π, we replace γ by the shortest path between A and B in $\mathcal{C}_i^\ell$ (resp. $\mathcal{C}_i^r$).

If $\mathcal{C}_i$ is concave, the straightening is defined as follows:
- If $\theta_r(A, B)$ and $\theta_\ell(A, B)$ are at least π, and even if $p_i \notin \gamma$, we replace γ by $[Ap_i] \cup [p_iB]$.
- If $\theta_r(A, B)$ (resp. $\theta_\ell(A, B)$) is strictly less than π, we replace γ by the shortest path between A and B in $\mathcal{C}_i^r$ (resp. $\mathcal{C}_i^\ell$).

In case $\gamma = c$ is strictly included in the interior of $\mathcal{C}_i$, then $\Phi_{loc}^i(c) = 0$, where 0 denotes an arbitrary constant curve based at a point p in $\mathcal{C}_i$.

We denote by Φ_{loc}^i, relative to a given star $\mathcal{C}_i$, the straightening process described above, applied in this star to each arc of a closed curve $c \in \Omega^{pl}$. Then Φ is defined as the concatenation $\Phi := \circ_{i=1}^n \Phi_{loc}^i$. Let us first show that Φ has values in Ω^{pl}, note that it suffices to prove it for Φ_{loc}^i. It is immediate that the image under Φ_{loc}^i is piecewise-linear. In order

to prove that the image is weakly simple, we look at the case of two arcs of the same closed curve c in a star, one delimited by two gates A and B, the other delimited by two gates A' and B'. As c belongs to Ω^{pl}, the two arcs do not cross, so they delimit a band in the star. If Φ^i_{loc} sends both arcs to the same side of p_i, then their images form two shortest paths in the same region and do not intersect. If Φ^i_{loc} sends the two arcs on opposite sides of p_i, a configuration where the two arcs cross twice is impossible because the angles $\theta_r(A, B)$ and $\theta_r(A', B')$ on the one hand, and $\theta_\ell(A, B)$ and $\theta_\ell(A', B')$ on the other hand are arranged in the same order.

If c is a quasigeodesic, each of its arcs possibly behaves in two ways in the star it crosses: either it reaches and leaves the vertex in a straight line from and up to the boundary of the star, forming on each side an angle at most π in the convex case, or at least π in the concave case. Or it connects its gates via a shortest path, entirely contained in the more acute of the two regions that it induces. In both cases, the previous process does not change its trajectory. So Φ fixes the quasigeodesics. Conversely, if c is not a quasigeodesic, then either it does not take a shortest path through a face or in neighborhood of a transverse intersection with an edge, either it forms on the passage of a vertex an angle greater than π on one side. This will be straightened when applying Φ^i_{loc} in a star containing that face, edge, or vertex in its interior, and therefore c is not a fixed point of Φ^i_{loc}. By construction, since Φ^i_{loc} does not increase lengths, we have that $L(\Phi(c)) \leq L(c)$. Let us show that if $L(\Phi(c)) = L(c)$, then c is a quasigeodesic. If an arc of c is not fixed by Φ^i_{loc} in a star, while remaining on the same side of the vertex, then it loses length, because there is uniqueness of the shortest path within a (left or right) region of a star. On the other hand, if Φ^i_{loc} passes an arc on the other side of the vertex (or pushes it against the vertex), it is because its length exceeds $L([Ap_i]) + L([p_iB])$. So the arc loses at least this excess in length. Finally, since some Φ^i_{loc} decreases the length of a non-quasigeodesic c, such a c cannot be a fixed point of Φ. ◀

In this proof, we could have taken the simpler choice of always replacing an arc in a star by a shortest path, irrespective of the angle at the vertex. The more delicate choice that is made here is tailored so as to be able to extend Φ to sweep-outs in Lemma 8.

The following property will be useful.

▶ **Lemma 6.** *For all $\varepsilon > 0$, there exists $\eta > 0$ such that for any curve $c \in \Omega^{pl}$ and for any i, if $L(c) - L(\Phi^i_{loc}(c)) < \eta$, then $d_\mathcal{H}(c, \Phi^i_{loc}(c)) < \varepsilon$, where $d_\mathcal{H}$ denotes the Hausdorff distance.*

The following lemma shows that applying Φ iteratively to a curve either makes the curve trivial in finite time, or converges to a quasigeodesic. Note that the lemma is not as obvious as it might seem as Φ is not continuous on Ω^{pl}.

▶ **Lemma 7.** *Let $c \in \Omega^{pl}$. We consider the sequence of iterates of Φ, i.e., $(\Phi^j(c))_j$. If this sequence does not reach 0 in finite time, then it admits a subsequence converging to a quasigeodesic (with respect to the uniform convergence metric).*

We now explain how to apply the disk flow to a monotone sweep-out, so that it extends the action on each of the fibers.

▶ **Lemma 8.** *There exists a map $\hat{\Phi} : \mathcal{B} \longrightarrow \mathcal{B}$ and a piecewise continuous injective map $\iota : [0, 1] \longrightarrow [0, 1]$, such that*

$$\forall s \in [0, 1], \hat{\Phi}(\beta)(\iota(s), \cdot) = \Phi(\beta(s, \cdot)).$$

The map ι induces a surjection f that maps $[0, 1]$ on to $[0, 1]$, which continually extends ι^{-1}, with the property that $L(\hat{\Phi}(\beta)(s, \cdot)) \leq L(\beta(f(s), \cdot))$, with equality if and only if $\beta(f(s), \cdot)$ is a quasigeodesic.

Proof. Let β be a sweep-out in $\mathcal{B}$. We explain how to apply a local step $\hat{\Phi}_{loc}^i$ of the curve-shortening process to β. Then, as before, we will define $\hat{\Phi}$ as the concatenation $\circ_{i=1}^n \hat{\Phi}_{loc}^i$.

Before analyzing the effect of Φ_{loc}^i on β, we apply an artificial thickening of β which fills its "problematic" portions on the boundary of each star and is defined as follows. We call the *bare boundary* of $\mathcal{C}_i$ the set of points of $\partial \mathcal{C}_i$ which are not the gates of any arc of a fiber of β crossing $\mathcal{C}_i$. Consider a connected component of the bare boundary of a certain star $\mathcal{C}_i$. It is fully contained in the image of at least[1] one fiber c of β that:

- either connects two gates which are neither a front gate nor an exit gate,
- or it connects a front or exit gate on one side only (see the green curve on Figure 2, top left),
- or it does not connect any gate (see the green curve on Figure 2, top right).

In all three cases, we can see that applying Φ_{loc}^i would induce a discontinuity around c. This is pictured in Figure 2, where one sees that the action of Φ_{loc}^i on the red curve and the green curve would be very different, despite them being arbitrarily close. We handle this discontinuity as follows. Case 1 will fit into the more general surgery described below, and thus is not addressed at this stage. In cases 2 and 3, the idea is to replace the parameter s of $c = \beta(s, \cdot)$ by a closed interval describing a collection of copies of c all identical (hence the artificial nature of this thickening), except that we drag artificially the position of the single extremal gate (case 2) or we add two new front/exit gates (case 3), one of which moves along $\partial \mathcal{C}_i$. In both cases, the new gates keep or gain an open character to the right or to the left. The aim of this operation is that the arcs of c between these new artificial gates will become straightened by Φ_{loc}^i, thus ensuring the continuity of Φ_{loc}^i at c (see Figure 2).
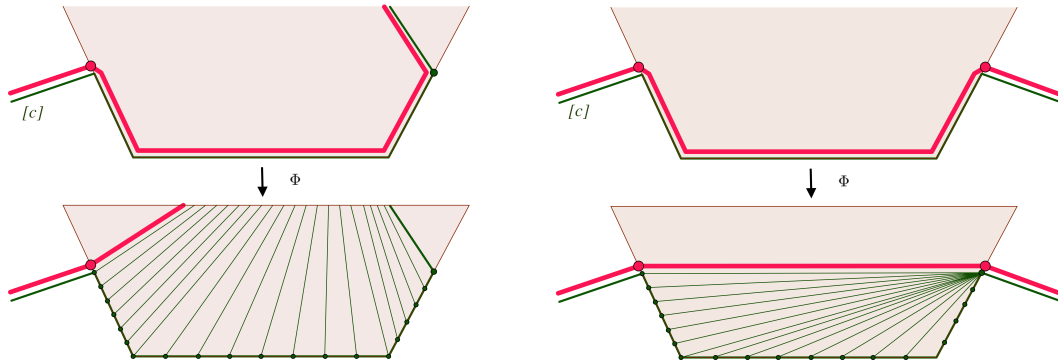
Figure 2 We artificially add gates on bare edges to obtain interpolating curves in their neighborhoods.

After this pre-processing, we consider the map $\beta' : [0, 1] \times \mathbb{S}^1 \longrightarrow S$ defined by:

$$\forall s \in [0, 1], \beta'(s, \cdot) = \Phi_{loc}^i(\beta(s, \cdot)).$$

The discontinuity of Φ_{loc}^i on arcs within the star $\mathcal{C}_i$ induces a finite number of tears in β'. These discontinuities occur around a convex vertex or when a fiber is tangent to the boundary of a star, and can be of the following four types.

---

[1] If there is an infinite number of them, they are parameterized in β by a closed interval. We then consider the representative closest to the interior of the star.

⬛ Disappearance of one or more gates far from the vertex, see Figure 3.
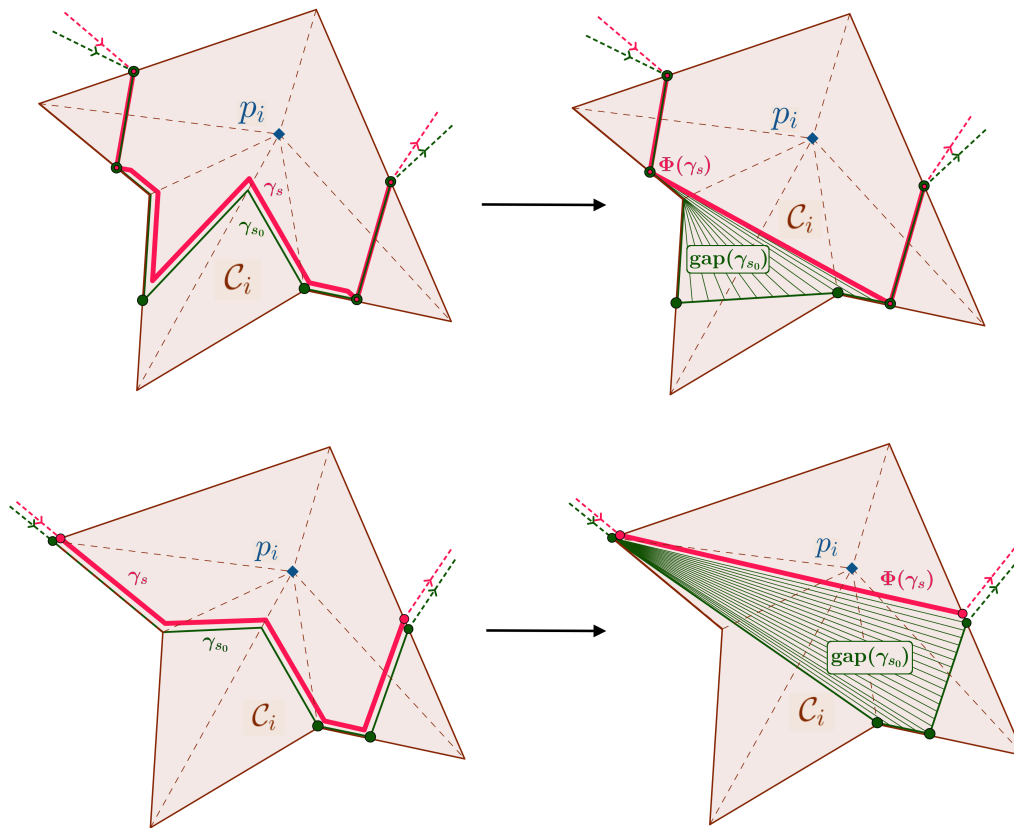
Figure 3 Disappearance of one or more gates far from the vertex: Two examples of interpolation. In the example in the bottom picture, the γ_{s_0} fiber has been added during the preprocessing and provided with an artificial gate. The missing part of the interpolation will be covered by the new gates induced in the preprocessing.

⬛ Double tear around a convex vertex, see Figure 4.
⬛ Single tear around a vertex, see Figure 5.
⬛ Disappearance of interior curves, see Figures 6 and 7.

In all four cases, the discontinuities are filled by (1) blowing up the parameter space around a discontinuity point s to an interval gap(s) and (2) adding interpolating curves in this interval gap(s), one of which is $\Phi^i_{loc}(\beta(s,\cdot))$. We refer to Athe full version for a precise description of these interpolating procedures. Such interpolating curves are pictured in green in Figure 3, 4, 5, 6 and 7 and they are obtained by shortcutting $\beta(s,\cdot)$ using shortest paths, therefore all of them have length bounded by that of $\beta(s,\cdot)$. We define the map ι as the one sending s to the parameter corresponding to the fiber $\Phi^i_{loc}(\beta(s,\cdot))$, while the surjection f maps the entire interval gap(s) to s (the maps ι and f are defined in the natural way outside of the discontinuities). Therefore we have defined a new map which we denote by $\hat{\Phi}^i_{loc}(\beta)$, whose parameter space is connected to that of β using the maps ι and f. As the $\hat{\Phi}^i_{loc}$ get composed to yield $\hat{\Phi}$, the maps ι and f are also composed in the natural way.
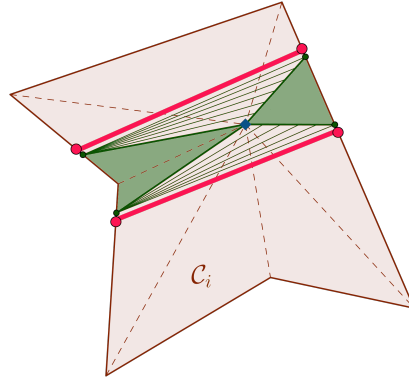
**Figure 4** Double tear around a convex vertex: Interpolation.

We argue that the resulting map is a monotone sweep-out. It starts and ends at trivial curves, and by constructions each fiber is piecewise-linear. Furthermore, the disks defined by the fibers are nested, since the effect of $\hat{\Phi}^i_{loc}$ is restricted to the star $\mathcal{C}_i$, where the nesting of disks that was present in β is preserved, as the interpolated curves are put inbetween their interpolation targets. Generically, points are covered by the new sweep-out exactly once (since all the fibers can be slightly perturbed to be disjoint), thus the topological degree is one. Finally, since all the interpolating curves have length at most that of a curve it interpolates from, we have the inequality $L(\hat{\Phi}(\beta)(s, \cdot)) \leq L(\beta(f(s), \cdot))$, with equality if and only if $\beta(f(s), \cdot)$ is a quasigeodesic.   ◀

▶ Remark. This proof showcases why our definition of quasigeodesic is the correct one for the disk flow to be appropriately defined on sweep-outs. If we had chosen more strict rules around convex vertices (for example only allowing curves with equal angles on both sides), we could have defined Φ in a more abrupt way by simply replacing arcs with shortest paths, thus ensuring that no arc through a vertex is fixed by the disk flow. However, this would have yielded tears around a convex vertex p in which our interpolating technique could not have worked, since no fiber of β' would be going through the vertex, and there would have been no way to add interpolating fibers of controlled length. In this sense, allowing for an angle at most π on both sides is the minimum angular spread allowing for the interpolation steps in the proof of Lemma 8 to work. For concave vertices, shortest paths between points on the boundary of a star $\mathcal{C}_i$ might require the whole spread of angles at least π on both sides, hence this choice of definition.

4   Existence of a simple closed quasigeodesic

We are now ready to prove Theorem 1. At this stage, our proof follows the same lines as that of Hass and Scott [16, Theorem 3.11].

Proof of Theorem 1. Let β be the monotone sweep-out of $\mathcal{B}$ of width at most M described by Lemma 4. We consider the sequence of sweep-outs $(\hat{\Phi}^j(\beta))_j$. For any $j \in \mathbb{N}$, the parameter space of $\hat{\Phi}^j(\beta)$ is the product of an interval $[0,1]$ by $\mathbb{S}^1$, the first factor being related to that of $\hat{\Phi}^{j-1}(\beta)$ via the surjection f_j of Lemma 8. Therefore, in order to track the history of a fiber in $\hat{\Phi}^j(\beta)$ under the action of $\hat{\Phi}$, we introduce the sequence of parameters $\mathcal{O}_j = (s_0, \ldots, s_j)$ such that for all k beetwen 0 and $j-1$: $s_k = f_k(s_{k+1})$. Each space of parameters describing
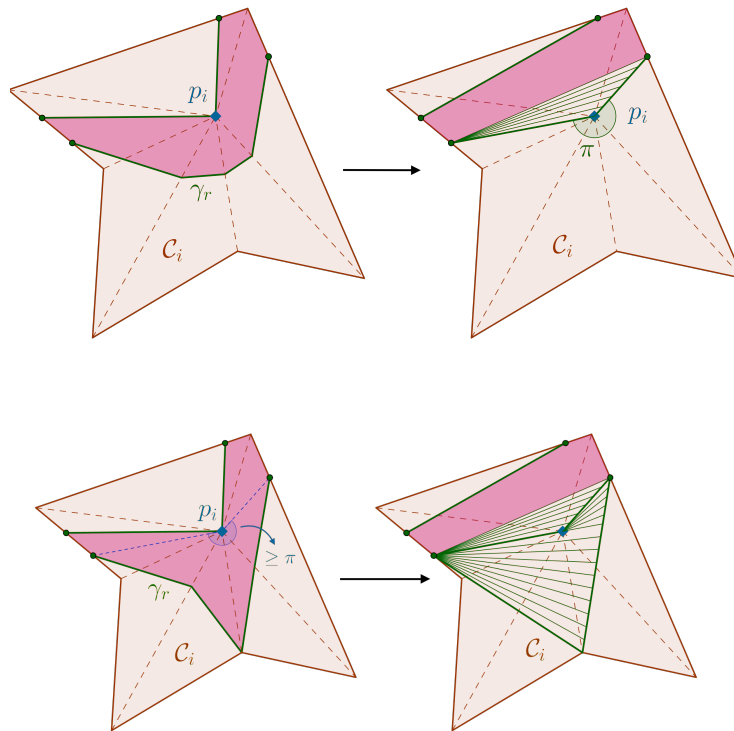
Figure 5 Single tear around a vertex: Two examples of interpolation.

$\mathcal{O}_j$ is homeomorphic to the interval $[0, 1]$ (via the trivial homeomorphism $(s_0, \ldots, s_j) \mapsto s_j$), and we consider the projective limit $\mathcal{I}$ of these intervals, which is thus also homeomorphic to an interval $[0, 1]$. An element of this projective limit therefore consists of an infinite sequence $\mathcal{O} = (s_0, s_1 \ldots)$ such that for all k, $s_k = f_k(s_{k+1})$.

Let $\mathcal{O} = (s_1, s_2 \ldots)$ be an element of $\mathcal{I}$, which thus corresponds to a family of curves $c_j := \hat{\Phi}^j(\beta)(s_j)$, and let us assume that all these curves are trivial for j bigger than some k. Then there is an open neighborhood of $\mathcal{O}$ for which this is also the case, as a curve becomes trivial under the action of some $\hat{\Phi}^i_{loc}$ if and only if it is fully contained in the interior of a star. Therefore, the set $\mathcal{V}^k \subset \mathcal{I}$ of sequences of curves for which the kth curve is not trivial is a closed subset of $\mathcal{I}$. Furthermore, it is not empty, as otherwise some intermediate sweep-out after $\hat{\Phi}^{k-1}(\beta)$ would consist of only curves contained in the interior of some star and thus would miss some point of the sphere S, in contradiction with the requirement that a sweep-out be of topological degree one. Finally, we have the natural inclusion $\mathcal{V}^{k+1} \subset \mathcal{V}^k$ since if $\hat{\Phi}^{k+1}(\beta)(s_{k+1})$ is not trivial, then this is also the case for $\hat{\Phi}^k(\beta)(s_k)$. We can thus consider the intersection $\cap_{k \in \mathbb{N}} \mathcal{V}^k$ which is an infinite intersection of nested closed non-empty subsets of $\mathcal{I}$ and is thus non-empty. An element in this intersection is a sequence $\mathcal{O}_\infty = (s_1, s_2 \ldots)$ such that none of the curves $c_n = \hat{\Phi}^n(\beta)(s_n)$ is trivial. As Ω^{pl} is compact, we can extract from this sequence of curves a convergent subsequence c_k, which converges to a curve c_∞. We claim that c_∞ is a weakly simple closed quasigeodesic of length at most M. The curve c_∞ is weakly simple because it is a limit of weakly simple curves. The bound on the length follows from the fact that by Lemma 4, the width of each of the sweep-outs $\hat{\Phi}^n(\beta)$ is at most M, and thus in particular c_∞ is a limit of curves of length at most M and thus has length at most M, since the length is a lower semi-continuous function on Ω.
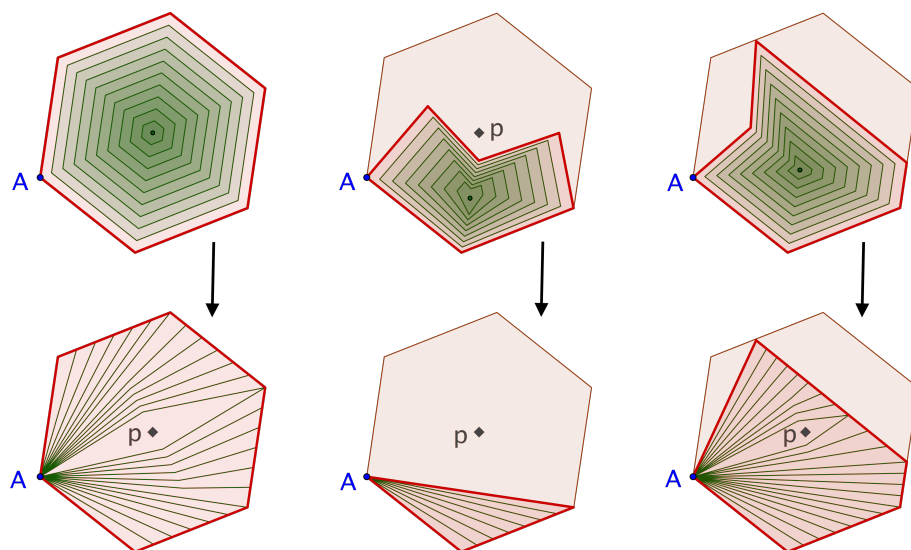
Figure 6 Disappearance of interior curves: Three examples of interpolation.
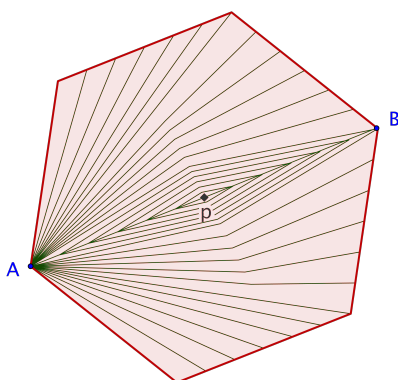
Figure 7 Interpolating to fill the last hole when an interior curve disappears.

Finally, in order to prove that c_∞ is a quasigeodesic, the argument is identical to the one in the proof of Lemma 7, to which we refer. If c_∞ is not a quasigeodesic, there is one point p in its image which is not locally quasigeodesic, i.e., there are two points p_1 and p_2 in a small neighborhood in c_∞ such that p_1, p and p_2 are not aligned, and if p is a vertex, the angle at p is disallowed by the curvature there. For k big enough, c_k will also have this property. Now, we consider a star $\mathcal{C}_i$ which contains p in its interior. Here observe that Lemma 6 is also valid under the action of $\hat{\Phi}$, i.e., for interpolating curves: indeed, those are displaced even less than under the action of Φ. Therefore, c_k will have moved very little when Φ_{loc}^i acts on it, and thus this action will diminish its length by a fixed quantity that can be lower bounded based on c_∞, which is impossible since the lengths of the c_k converge.   ◄

Our techniques only guarantee the existence of a weakly simple closed quasigeodesic of length at most M. In contrast, in the convex case, Pogorelov [19] proved the existence of a simple closed quasigeodesic (where the degenerate case of two vertices of curvature at least π connected twice by a curve is allowed). The proof of Pogorelov works by approximating

a convex polyhedron by smooth surfaces, appealing to the Lyusternik-Schnirrelmann on the smooth surfaces to find simple closed geodesics, taking the limit of such simple closed geodesics and arguing that (1) the limit is a quasigeodesic and (2) it is simple. We argue that the same technique proves the existence of a simple closed quasigeodesic of length at most $M + \varepsilon$, for an arbitrarily small $\varepsilon > 0$. Indeed, the sweep-out that we describe on S naturally induces sweep-outs of width at most $M + \varepsilon$ on the approximating smooth surfaces that are close enough, and thus the first simple closed geodesic output by the Lyusternik-Schnirrelmann theorem in each of these surfaces has length at most $M + \varepsilon$. Taking the limit of those yields a simple closed quasigeodesic of length at most $M + \varepsilon$. We will use this result at the end of the next section.

5   An algorithm to compute a weakly simple closed quasigeodesic

In this section, we leverage the existence of a weakly simple closed quasigeodesic of length at most M proved in Theorem 1 in order to design an algorithm to find it.

Let S be a polyhedral sphere and denote by $\mathcal{E} = \{p_1, \ldots, p_n, e_1, \ldots, e_m\}$ be the set of vertices and open edges of S. To a closed curve $c : \mathbb{S}^1 \longrightarrow S$, we associate the cyclic word $\mathcal{E}(c)$ whose successive letters are the elements of $\mathcal{E}$ met by $c(t)$ as t moves around $\mathbb{S}^1$ (note that an edge can be either crossed or followed). Given a bound on the length of c, we want to derive a bound on the combinatorics of c, i.e., a bound on the length of $\mathcal{E}(c)$. This is hopeless without any assumption, as a curve spiraling around a vertex for an arbitrarily long time showcases. But when c is a weakly simple quasigeodesic, we can obtain such a bound. Indeed, our first observation is that a weakly simple quasigeodesic never spirals around a vertex.

▶ **Lemma 9.** *Let γ be a weakly simple closed quasigeodesic and $\mathcal{C}_i$ be the open star of a vertex v_i of degree d_i. Then for any connected component α of $\gamma \cap \mathcal{C}_i$, the number of intersections of α with edges and vertices of $\mathcal{C}_i$ is at most d_i.*

Proof. If α passes through the vertex v_i, then it exits on both sides tracing a straight line in one of the triangles of $\mathcal{C}_i$. This straight-line reaches directly the opposite edge of the triangle, therefore in this case the number of intersections of α with edges and vertices of $\mathcal{C}_i$ is at most two.

If α does not pass through the vertex v_i, then let us denote by e the first edge adjacent to v_i that it crosses. Note that within a triangle of $\mathcal{C}_i$, by quasigeodesicity, α enters from one edge and does not backtrack, i.e., it escapes from another edge. Therefore, either α escapes from $\mathcal{C}_i$ before crossing e again, in this case it crosses at most d_i edges, or it crosses e again. In the latter case, up to reversing orientation of α we can assume that the second crossing point is closer to v_i than the first crossing point. Tracing α after the second crossing point, we see that in each triangle that it enters, it cannot escape $\mathcal{C}_i$ since, by weak simplicity, it cannot cross the previous edge that it traced, and is thus forced to continue spiraling around v_i indefinitely. This contradicts the assumption that γ is closed, finishing the proof. ◀

The following geometric lemma will come handy to bound the combinatorics of a closed simple quasigeodesic.

▶ **Lemma 10.** *Let Q be a Euclidean quadrilateral consisting of two Euclidean triangles glued along an edge. Then the distance between two opposite sides of Q is lower bounded by the smallest altitude of the two triangles.*

Combining Lemmas 9 and 10 yields the following proposition.

▶ **Proposition 11.** *Let S be a polyhedral sphere, let M denote the sum of the edge-lengths of the triangles of S, let h denote the smallest altitude of the triangles of S, and let d be the maximum degree of a vertex in S. Then there exists a weakly simple closed quasigeodesic on S such that the length of $\mathcal{E}(\gamma)$ is bounded by:*

$$\eta_\gamma = \lceil (d+1)M/h \rceil$$

We now have all the tools to prove Theorem 2.

Proof of Theorem 2. Let γ be a weakly simple closed quasigeodesic whose combinatorial complexity is controlled by η_γ as specified by Proposition 11. First, we observe that we can assume that this quasigeodesic meets a vertex (see full version).
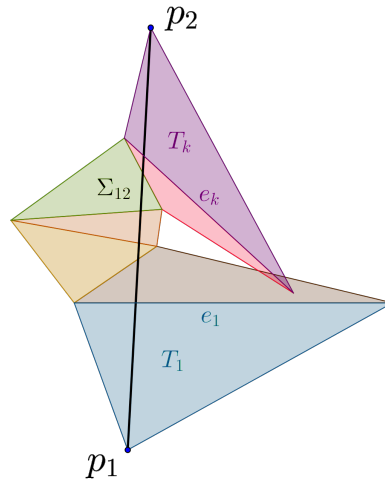
■ **Figure 8** Unfolding a (tentative) quasigeodesic along the set of edges that it crosses.

Then, we guess the cyclic word w of size at most η_γ describing the combinatorics of γ, a weakly simple closed quasigeodesic going through at least one vertex. For each subword $p_1 e_1 \dots e_k p_2$ between two consecutive vertices, if e_1 is adjacent to p_1, we simply check that the next letter is the other endpoint of p_1. Otherwise, we first check that successive letters of that word are adjacent to a common triangle. Then we compute a local unfolding of the polyhedral sphere along the edges $e_1, e_2 \dots e_k$, i.e., we first place the triangle T_1 spanned by p_1 and e_1, to which we attach along e_1 the triangle T_2 spanned by e_2 and e_3, and so on until we reach the last triangle T_k spanned by e_k and p_2. Now, in this unfolded picture, we trace the straight line Σ_{12} between p_1 and p_2. There remains to check that the combinatorics of this straight line match those of the guessed word: in the first and last triangles, we check that Σ_{12} exits via or follows e_1 (or e_k), and in each other triangle T_i it suffices to check that the three vertices of T_i are on the sides of Σ_{12} prescribed by the edges e_i and e_{i+1} (i.e., if $e_i = ab$ and $e_{i+1} = bc$, then a and c should be one side of Σ_{12} while b should be on the other side). Then, we check that the angle between each pair $\Sigma_{i,i+1}, \Sigma_{i+1,i+2}$ is within the rules specified by the curvature at the vertex p_{i+1}. Finally, we check that this curve is weakly simple, for example via known algorithms [1, 10] or by brute-forcing in exponential time the choice of on which side two overlapping segments can be desingularized. If all the checks are positive, we have found the unique closed quasigeodesic matching the combinatorics of the word w, which is thus weakly simple. ◀

Finally, let us discuss how to find a simple closed quasigeodesic of bounded length in the case of a convex polyhedron. Following the discussion at the end of Section 4, Pogorelov's theorem implies that there exists a simple closed quasigeodesic of length at most $M + \varepsilon$, for an arbitrarily small ε, and allowing as a "simple" closed quasigeodesic the degenerate case of a curve connecting twice two vertices of curvature at least π. This degenerate case is a weakly simple curve that will be found by our algorithm. For the non-degenerate case, the arguments of Proposition 11 apply verbatim to provide a bound on the combinatorics of some simple closed quasigeodesic γ. If this quasigeodesic goes through at least one vertex, the algorithm described just above finds it, and it is immediate to check that it is simple. If not, we can push it as in the proof of Theorem 2 to a weakly simple closed geodesic that goes through a vertex, and it will stay simple until it hits that vertex, where it will form an angle exactly π in the direction where it came from. Since the total angle at each vertex is at most 2π, this implies that this curve is either degenerate or simple, and in both cases it will be found by our algorithm.

―――― **References** ――――

1   Hugo A Akitaya, Greg Aloupis, Jeff Erickson, and Csaba D Tóth. Recognizing weakly simple polygons. *Discrete & Computational Geometry*, 58(4):785–821, 2017. `doi:10.1007/s00454-017-9918-3`.

2   Alexandr D Alexandrov. *Convex polyhedra*. Springer Science & Business Media, 2005.

3   Werner Ballmann. Der Satz von Lusternik und Schnirelmann. *Bonner Math. Schriften*, 102:1–25, 1978.

4   Werner Ballmann, Gudlaugur Thorbergsson, and Wolfgang Ziller. On the existence of short closed geodesics and their stability properties. In *Seminar On Minimal Submanifolds.(AM-103), Volume 103*, pages 53–64. Princeton University Press, 1983.

5   George David Birkhoff. *Dynamical systems*, volume 9 of *Colloquium Publications*. American Mathematical Soc., 1927. `doi:10.1016/B978-044450871-3/50149-2`.

6   Heinz Bruggesser and Peter Mani. Shellable decompositions of cells and spheres. *Mathematica Scandinavica*, 29(2):197–205, 1971. `doi:10.7146/math.scand.a-11045`.

7   Dmitri Burago, Yuri Burago, and Sergei Ivanov. *A course in metric geometry*, volume 33. American Mathematical Society, 2001.

8   Yuriy Dmitrievich Burago and Viktor Abramovich Zalgaller. Isometric piecewise-linear embeddings of two-dimensional manifolds with a polyhedral metric into $\mathbb{R}^3$. *Algebra i analiz*, 7(3):76–95, 1995.

9   Erin Wolf Chambers, Gregory R Chambers, Arnaud de Mesmay, Tim Ophelders, and Regina Rotman. Constructing monotone homotopies and sweepouts. *Journal of Differential Geometry*, 119(3):383–401, 2021. `doi:10.4310/jdg/1635368350`.

10   Hsien-Chih Chang, Jeff Erickson, and Chao Xu. Detecting weakly simple polygons. In *Proceedings of the twenty-sixth annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1655–1670. SIAM, 2014. `doi:10.1137/1.9781611973730.110`.

11   Jean Chartier and Arnaud de Mesmay. Finding weakly simple closed quasigeodesics on polyhedral spheres, 2022. `arXiv:2203.05853`.

12   Erik D Demaine, Adam C Hesterberg, and Jason S Ku. Finding closed quasigeodesics on convex polyhedra. In *36th International Symposium on Computational Geometry (SoCG 2020)*, pages 33:1–33:13. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020. `doi:10.1007/b137434`.

13   Erik D Demaine and Joseph O'Rourke. *Geometric folding algorithms: linkages, origami, polyhedra*. Cambridge University Press, 2007. `doi:10.1017/CBO9780511735172`.

14   Jeff Erickson and Amir Nayyeri. Tracing compressed curves in triangulated surfaces. *Discrete & Computational Geometry*, 49(4):823–863, 2013. `doi:10.1007/s00454-013-9515-z`.

15   Matthew A Grayson. Shortening embedded curves. *Annals of Mathematics*, 129(1):71–111, 1989. `doi:10.2307/1971486`.

16   Joel Hass and Peter Scott. Shortening curves on surfaces. *Topology*, 33(1):25–43, 1994. `doi:10.1016/0040-9383(94)90033-7`.

17   Daniel Kane, Gregory N Price, and Erik D Demaine. A pseudopolynomial algorithm for Alexandrov's theorem. In *Workshop on Algorithms and Data Structures*, pages 435–446. Springer, 2009. `doi:10.1007/978-3-642-03367-4_38`.

18   L Lyusternik and Lev Schnirelmann. Sur le problème de trois géodésiques fermées sur les surfaces de genre 0. *CR Acad. Sci. Paris*, 189(269):271, 1929.

19   Aleksei Vasil'evich Pogorelov. Quasi-geodesic lines on a convex surface. *Matematicheskii Sbornik*, 67(2):275–306, 1949. English translation in *American Mathematical Society Translations* 74, 1952.

20   Henri Poincaré. Sur les lignes géodésiques des surfaces convexes. *Transactions of the American Mathematical Society*, 6(3):237–274, 1905.

21   Nicholas Sharp, Yousuf Soliman, and Keenan Crane. Navigating intrinsic triangulations. *ACM Transactions on Graphics (TOG)*, 38(4):1–16, 2019. `doi:10.1145/3306346.3322979`.

22   Günter M Ziegler. *Lectures on polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer Science & Business Media, 2012. `doi:10.1007/978-1-4613-8431-1`.