Finding Closed Quasigeodesics on Convex Polyhedra

Erik D. Demaine
Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge, MA, USA
edemaine@mit.edu

Adam C. Hesterberg
Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge, MA, USA
achester@mit.edu

Jason S. Ku
Department of Electrical Engineering and Computer Science, MIT, Cambridge, MA, USA
jasonku@mit.edu

—— **Abstract** ——

A **closed quasigeodesic** is a closed loop on the surface of a polyhedron with at most 180° of surface on both sides at all points; such loops can be locally unfolded straight. In 1949, Pogorelov proved that every convex polyhedron has at least three (non-self-intersecting) closed quasigeodesics, but the proof relies on a nonconstructive topological argument. We present the first finite algorithm to find a closed quasigeodesic on a given convex polyhedron, which is the first positive progress on a 1990 open problem by O'Rourke and Wyman. The algorithm's running time is pseudopolynomial, namely $O\left(\frac{n^2}{\varepsilon^2}\frac{L}{\ell}b\right)$ time, where ε is the minimum curvature of a vertex, L is the length of the longest edge, ℓ is the smallest distance within a face between a vertex and a nonincident edge (minimum feature size of any face), and b is the maximum number of bits of an integer in a constant-size radical expression of a real number representing the polyhedron. We take special care in the model of computation and needed precision, showing that we can achieve the stated running time on a pointer machine supporting constant-time w-bit arithmetic operations where $w = \Omega(\lg b)$.

1 Introduction

A **geodesic** on a surface is a path that is locally shortest at every point, i.e., cannot be made shorter by modifying the path in a small neighborhood. A **closed geodesic** on a surface is a loop (closed curve) with the same property; notably, the locally shortest property must hold at all points, including the "wrap around" point where the curve meets itself. In 1905, Poincaré [20] conjectured that every convex surface has a non-self-intersecting closed geodesic.[1] In 1927, Birkhoff [5] proved this result, even in higher dimensions (for any smooth metric on the n-sphere). In 1929, Lyusternik and Schnirelmann [17] claimed that every smooth surface of genus 0 in fact has at least *three* non-self-intersecting closed geodesics. Their argument "contains some gaps" [2], filled in later by Ballmann [1].

---

[1] Non-self-intersecting (quasi)geodesics are often called *simple* (quasi)geodesics in the literature; we avoid this term to avoid ambiguity with other notions of "simple".

■ **Figure 1** At a vertex of curvature κ, there is a κ-size interval of angles in which a segment of a quasigeodesic can be extended: the segment of geodesic starting on the left can continue straight in either of the pictured unfoldings, or any of the intermediate unfoldings in which the right pentagon touches only at a vertex.

For non-smooth surfaces (such as polyhedra), an analog of a geodesic is a ***quasigeodesic*** – a path with $\leq 180°$ of surface on both sides locally at every point along the path. Equivalently, a quasigeodesic can be locally unfolded to a straight line: on a face, a quasigeodesic is a straight line; at an edge, a quasigeodesic is a straight line after the faces meeting at that edge are unfolded (developed) flat at that edge; and at a vertex of curvature κ (that is, a vertex whose sum of incident face angles is $360° - \kappa$), a quasigeodesic entering the vertex at a given angle can exit it anywhere in an angular interval of length κ, as in Figure 1. Analogously, a ***closed quasigeodesic*** is a loop which is quasigeodesic. In 1949, Pogorelov [19] proved that every convex surface has at least three non-self-intersecting closed quasigeodesics, by applying the theory of quasigeodesics on smooth surfaces to smooth approximations of arbitrary convex surfaces and taking limits.

The existence proof of three closed quasigeodesics is nonconstructive, because the smooth argument uses a nonconstructive topological argument (a homotopy version of the intermediate value theorem).[2] In 1990, Joseph O'Rourke and Stacia Wyman posed the problem of finding a polynomial-time algorithm to find any closed quasigeodesic on a given convex polyhedron (aiming in particular for a non-self-intersecting closed quasigeodesic) [18]. This open problem was stated during the open problem session at SoCG 2002 (by O'Rourke) and finally appeared in print in 2007 [9, Open Problem 24.24]. Two negative results mentioned in [9] are that an n-vertex polyhedron can have $2^{\Omega(n)}$ non-self-intersecting closed quasigeodesics (an unpublished result by Aronov and O'Rourke) and that, for any k, there is a convex polyhedron whose shortest closed geodesic is not composed of k shortest paths (an unpublished result from the discussion at SoCG 2002).

Even a finite algorithm is not known or obvious. One approach is to argue that there is a closed quasigeodesic consisting of $O(n)$ (or any function $s(n)$) segments on faces. It seems plausible that the "short" closed quasigeodesics from the nonconstructive proofs satisfy

---

[2] A proof sketch for the existence of one closed geodesic on a smooth convex surface is as follows. By homotopy, there is a transformation of a small clockwise loop into its (counterclockwise) reversal that avoids self-intersection throughout. Consider the transformation that minimizes the maximum arclength of any loop during the transformation. By local cut-and-paste arguments, the maximum-arclength intermediate loop is in fact a closed geodesic. The same argument can be made for the nonsmooth case.

this property, but as far as we know the only proved property about them is that they are non-self-intersecting, which does not obviously suffice. (A quasigeodesic could plausibly wind many times around a curvature-bisecting loop, like the equator of a cube, somehow turn around, and symmetrically unwind, all without collisions.) If true, there are $O(n)^{s(n)}$ combinatorial types of quasigeodesics to consider, and each can be checked via the existential theory of the reals (in exponential time), resulting in an exponential-time algorithm. But we do not know how to bound $s(n)$; even the results of this paper give no upper bound on the number of segments constituting a closed quasigeodesic. In general, polyhedra such as isosceles tetrahedra have arbitrarily long non-self-intersecting closed geodesics (and even infinitely long non-self-intersecting geodesics) [14], so the only hope is to find an upper bound $s(n)$ on some (fewest-edge) closed quasigeodesic.

1.1 Our Results

We develop an algorithm that finds at least one[3] closed quasigeodesic on a given convex polyhedron in $O\left(\frac{n^2}{\varepsilon^2}\frac{L}{\ell}\right)$ real operations, where n is the number of vertices of the polyhedron, ε is the smallest curvature at a vertex, L is the length of the longest edge, and ℓ is the smallest distance within a face between a vertex and a nonincident edge (minimum feature size of any face). In the model described below in Section 1.2, these real operations take $O\left(\frac{n^2}{\varepsilon^2}\frac{L}{\ell}b\right)$ time if the input numbers are constant-size radical expressions over b-bit integers.

This running time is pseudopolynomial, so this does not yet resolve the open problem of a polynomial-time algorithm. The closed quasigeodesic output by our algorithm may be self-intersecting, even though a non-self-intersecting closed quasigeodesic is guaranteed to exist. Furthermore, the quasigeodesic path is output implicitly (in a format detailed below), as we lack a bound on the number $s(n)$ of needed segments. In Section 3, we discuss some of the difficulties involved in resolving either of these issues.

1.2 Models of Computation

Our results hold in several standard models of computation, which we pay careful attention to. While much work has been done on computational geometry in bounded-precision computational models [7, 8, 23], we are not aware of a single reference that describes all the relevant models, or with models that can handle representing polyhedra and geodesic paths. Thus we detail the possible model choices, distinguishing four aspects of the model:

1. **Combinatorial model of computation.** In combinatorial algorithms and data structures (without real numbers), there are two popular models of computation:

   a. **Pointer machine** [3, 10, 13]: Memory is decomposed into m records, each with a constant number of fields. Each field can store a w-bit integer or a pointer to another record. One record represents the machine's registers, and in constant time, the machine can read and write fields within constant pointer distance from that record, and create and/or destroy such a record.

---

[3] Our algorithm may in fact produce a list of closed quasigeodesics, but there are some closed quasigeodesics that it cannot find, including closed geodesics (not passing through a vertex) and possibly all non-self-intersecting closed quasigeodesics.

b. **Word RAM** [11]: Memory is an array of m words, each of which can store a w-bit integer. The first $O(1)$ words represent the machine's registers. In constant time, the machine can modify the register words or any word whose array index is given by a register word. The word RAM can simulate the pointer machine.

In both cases, the machine can also do basic w-bit integer arithmetic ($+, -, \times, \div,$ mod, AND, OR, NOT, $<, >, =$) in constant time, and we assume that $w = \Omega(\lg n)$ (the transdichotomous assumption, named for how it bridges the problem and machine [11]). Our algorithm works in the pointer machine, and thus also in the word RAM.

2. **Real model of computation.** To deal with geometry (e.g., to represent the input), we need to handle some form of real numbers. We define three models of increasing generality that represent numbers in some binary format.

a. **Integers** (encoded in binary): Can be added, subtracted, multiplied, and compared in $O(b)$ time, where b is the total number of bits in the operands. Here we make the transdichotomous assumption that $w = \Omega(\lg b)$, in which case integer multiplication can be done in $O(b)$ time [16]. (Without this assumption, integer multiplication requires $O(b \log b)$ bit operations [12], so we would just gain a log factor in our running times.)

b. **Rationals** (encoded as two integers, numerator and denominator, in binary): Same performance as integers, plus real division in $O(b)$ time.

c. **Constant-size radical expressions over integers** (encoded as a constant-size expression tree with operators $+, -, \times, \div, \sqrt[k]{}$ and integer leaves): By known results in root separation bounds [6], if b is the total number of bits in the integer leaves, then the first $O(b)$ bits can be computed in $O(b)$ time; and if the expression is nonzero, some of these $O(b)$ bits will be nonzero, enabling exact comparison with 0. Thus, such numbers can be compared in $O(b)$ time, but when combining them arithmetically, we need to take care that the expression never grows beyond constant size. Also, we can compute the floor of such a number in $O(b)$ time.

Our algorithm works in the last model, and therefore supports inputs in any of the three models. Past work [7, 8] assumes $O(w)$-bit integer or rational inputs, which is less suitable for inputs of polyhedra, as described below.

These models contrast the real RAM model [22], where the inputs are black-box real numbers supporting radical operations $+, -, \times, \div, \sqrt[k]{}$ and comparisons in constant time. While standard in computational geometry, this model is not very realistic for a digital computer, because it does not bound the required precision, which can grow without bound. For example, the real RAM model crucially does not support converting black-box real numbers into integers (e.g., via the floor function), or else one can solve PSPACE [21] and #SAT [4] in polynomial time. Our algorithm actually needs to use the floor function, so it does not work on the "reasonable" real RAM model which lacks this operation.

3. **Polyhedron format.** The combinatorial structure of the polyhedron can be encoded as a primal or dual graph, as usual, but which real numbers should represent the geometry? Because the quasigeodesic problem is about the intrinsic geometry of the surface of a polyhedron, the input geometry can be naturally represented intrinsically as well as extrinsically, leading to three natural representations:

a. **Extrinsic coordinates:** 3D coordinates for each vertex.

b. **Intrinsic coordinates:** For each face, for some isometric embedding of the face into 2D, the 2D coordinates of each vertex of the embedded face.

c. **Intrinsic lengths:** For each face, the lengths of the edges. This representation assumes the faces have been combinatorially triangulated (so some edges may be flat).

In our real-number model of constant-size radical expressions over integers, extrinsic coordinates can be converted into intrinsic coordinates which can be converted from/to intrinsic lengths. Indeed, this feature is one of our motivations for this real-number model. (The reverse direction, from intrinsic to extrinsic, is more difficult, as it involves solving the Alexandrov problem [15].)

Our algorithm works in any of these input models.

4. **Output format.** Because we do not know any bounds on the number of segments (faces) in a closed quasigeodesic, we need to allow an implicit representation of the output. Specifically, we allow a quasigeodesic to be specified by a sequence of commands of the following form:

follow a path R from vertex u to vertex v, traversing through some prefix of faces in the periodic sequence

$$f_1, f_2, \ldots, f_m; f_1, f_2, \ldots, f_m; \ldots.$$

We cannot specify the number of faces in the prefix, nor the length $\ell(R)$ of the path, but we can compute $\ell(R)$ with any desired precision: specifically, we can compute $\Omega(k)$ high-order bits of $\ell(R)$ in $O(k)$ time. We also cannot specify the direction that the path leaves u with exact precision, but we can compute $\Omega(k)$ high-order bits of the coordinates for a point p such that the direction of the path leaving u is toward p, again in $O(k)$ time. None of these quantities can be specified exactly, because the number of needed bits may be arbitrarily large, again because the path may visit arbitrarily many faces. But we can guarantee that such a path exists.

2 Algorithm

In this section, we give an algorithm to find a closed quasigeodesic on the surface of a convex polyhedron P. First, a bit of terminology: we define a *(quasi)geodesic ray/segment* to be a one/two-ended path that is (quasi)geodesic.

2.1 Outline

The idea of the algorithm is roughly as follows: first, we define a directed graph for which each node[4] is a pair $(V, [\varphi_1, \varphi_2])$ of a vertex V of P and a small interval of directions at it, with an edge from one such node, (U, I), to another, (V, J), if a geodesic ray starting at the polyhedron vertex U and somewhere in the interval of directions I can reach V and continue quasigeodesically everywhere in J[5]. We show how to calculate at least one out-edge from every node of that graph, so we can start anywhere and follow edges until hitting a node twice, giving a closed quasigeodesic.

---

[4] We use the word "node" and lower-case letters for vertices of the graph to distinguish them from vertices of a polyhedron, for which we use capital letters and the word "vertex".

[5] Since we consider only geodesic rays that can continue quasigeodesically *everywhere* in J, there are some closed quasigeodesics that we cannot find: those that leave a polyhedron vertex in a direction in an interval J for which some directions are not quasigeodesic continuations. In particular, this algorithm is unlikely to find closed quasigeodesics that turn maximally at a polyhedron vertex.

■ **Figure 2** A segment of a geodesic is a straight line in the unfolding of the sequence of faces through which it passes, as in this unfolding of a regular dodecahedron.

The key part of this algorithm is to calculate, given a polyhedron vertex U and a range of directions as above, another vertex V that can be reached starting from that vertex and in that range of directions, even though reaching V may require crossing superpolynomially many faces. First we prove some lemmas toward that goal.

▶ **Definition 2.1.** *If X is a point on the surface of a polyhedron, φ is a direction at X, and $d > 0$, then $R(X, \varphi, d)$ is the geodesic segment starting at X in the direction φ and continuing for a distance d or until it hits a polyhedron vertex, whichever comes first.[6] We allow $d = \infty$; in that case, $R(X, \varphi, d)$ is a geodesic ray.*

▶ **Definition 2.2.** *If $R(X, \varphi, d)$ is a geodesic segment or ray, the **face sequence** $F(R(X, \varphi, d))$ is the (possibly infinite) sequence of faces that $R(X, \varphi, d)$ visits.*

▶ **Lemma 2.3.** *If $R_1 = R(X, \varphi_1, \infty)$ and $R_2 = R(X, \varphi_2, \infty)$ are two geodesic rays from a common starting point X with an angle between them of $\theta \in (0, \pi)$, the face sequences $F(R_1)$ and $F(R_2)$ are distinct, and the first difference between them occurs at most one face after a geodesic distance of $O(L/\theta)$.*

Proof. Given a (prefix of) $F(R_i)$, the segment of R_i on it is a straight line, so while $F(R_1) = F(R_2)$, the two geodesics R_1 and R_2 form a wedge in a common unfolding, as in Figure 2. The distance between the points on the rays at distance d from X is $2d \sin \frac{\theta}{2} > d\theta/\pi$ (since $\frac{\theta}{2} < \frac{\pi}{2}$), so at a distance of $O(L/\theta)$, that distance is at least L. So either $F(R_1)$ and $F(R_2)$ differ before then, or the next edge that R_1 and R_2 cross is a different edge, in which case $F(R_1)$ and $F(R_2)$ differ in the next face, as claimed. ◀

If we had defined L analogously to ℓ as not just the length of the longest edge but the greatest distance within a face between a polyhedron vertex and an edge not containing it, we could remove the "at most one face after" condition from Lemma 2.3.

---

[6] This definition is purely geometric; we reserve calculating these paths for Lemma 2.4.

Figure 3 Even a short geodesic path between two vertices u and v may cross many edges. Equally colored faces represent copies of the same face being visited multiple times.

Figure 4 If a geodesic path encounters the same edge twice in nearly the same place and nearly the same direction, as is the case for the thick quasigeodesic path through the center of this figure if every fourth triangle is the same face, it may pass the same sequence of faces in the same order a superpolynomial number of times. Equally colored faces represent copies of the same face being visited multiple times.

2.2 Extending Quasigeodesic Rays

Although Lemma 2.3 gives a bound on the geodesic distance to the first difference in the face sequences (or one face before it), this gives no bound on the number of faces traversed before that difference, which might be large if the two paths come very close to a polyhedron vertex of high curvature, as in Figure 3, or repeat the same sequence of edges many times, as in Figure 4.

Nonetheless, in both of these cases, we can describe a geodesic ray's path efficiently:

▶ **Lemma 2.4.** *Let $R = (X, \varphi, d)$ be a geodesic segment with $d < \ell$. In $O(nb)$ time, we can calculate $F(R)$, expressed as a sequence S_1 of $O(n)$ faces, followed by another sequence S_2 of $O(n)$ faces and a distance over which R visits the faces of S_2 periodically[7]. Also, we can calculate the face, location in the face, and direction of R at its far endpoint (the one other than X).*

Proof. First, we prove the geometric fact (without calculating anything) that R is periodic from the first time it reenters any already-visited face. Second, we calculate, in $O(nb)$ time, the path of R through the non-periodic part, possibly detecting that R hits a vertex. Third, we calculate the path of R through the periodic part, in two cases: either R reenters each face at the same angle as its first entry, or not.

First, we claim that R is periodic from the first time it reenters a vertex: that is, if R enters a face f on an edge e_1 and exits[8] at a point P_2 on an edge e_2, then we claim that *every* time R enters f by e_1, it must exit f by e_2, and not any other edge e_3[9]. It must exit by a different edge from the edge e_1 by which it entered, so suppose for contradiction that in some visit to f, it enters at a point P_1 on the edge e_1 and exits at a point P_3 on another edge e_3, as shown in Figure 5. If any two of e_1, e_2, and e_3 are nonincident, then R has gone from a point on one edge to a point on a nonincident edge. By the definition of ℓ, R cannot do so without traveling a distance at least ℓ, farther than the conditions under which this lemma applies. Otherwise, e_1, e_2, and e_3 are the three edges of a triangular face, and the total geodesic distance is at least $d(P_1, P_2) + d(P_1, P_3)$. Consider the reflection e_4 of e_3 across e_2 and the reflected point P_4 on e_4. The path from P_4 to P_2 via P_1 is at least the distance from P_2 to P_4, which is at least the shortest distance from a point on e_4 to a point on e_2, which is attained at an endpoint of at least one of e_2 and e_4, say an endpoint of e_4. The path making that shortest distance (shown in gray) goes through e_1, so R travels at least the distance from e_1 to the opposite vertex, which is at least ℓ, farther than the conditions under which this lemma applies. Hence each edge crossed determines the next edge crossed, so $F(R)$ is periodic after crossing each edge at most once. Also, there are only $O(n)$ edges, so after crossing at most $O(n)$ edges, $F(R)$ repeats periodically with period $O(n)$.

Second, in total time $O(nb)$, we calculate the path of R before it repeats periodically in each face f it enters. Assume we start with an intrinsic representation of the polyhedron, with an isometric embedding of each face. We will represent the direction of a ray in a face by a pair of points, both with $O(b)$ bits, on the ray in the local coordinate system of that face; the points may or may not themselves be in the face.

Suppose that R enters f on an edge e' from a face f'. The isometry that takes the instance of e' in the embedding of f' to the instance of e' in the embedding of f is a linear transformation whose coefficients are $O(b)$ bits, which we can apply in $O(b)$ time to the pair of points representing R in f' to get a pair of points (x_0, y_0) and (x_1, y_1) representing R in f, again rounded to $O(b)$ bits.

For each edge e of f with endpoints (a, b) and (c, d), the intersection of the extension of R, which has equation $(x - x_1)(y_0 - y_1) = (x_0 - x_1)(y - y_1)$, and the extension of e, which has equation $(x - c)(b - d) = (a - c)(y - d)$, is a point (x, y) where each of x and y is a constant-depth arithmetic expressions in x_0, y_0, x_1, y_1, a, b, c, and d, so we can compute it in $O(b)$ time. Then we can check whether x is between a and c (or y is between b and d);

---

[7] The length of the sequence of faces may be too large to even write down the number of repetitions.
[8] If R hits a vertex of that polyhedron face f, we say that it exits on each of the two edges of f containing that vertex.
[9] In particular, R cannot exit by a vertex in any visit to f after the first.

Figure 5 If a geodesic visits three edges of the same face, the total distance traveled is at least ℓ.

Figure 6 When a quasigeodesic path passes through the same sequence of faces several times, the unfolding of the faces it passes through repeats regularly.

having $O(b)$ bits of each is enough to do so by Section 1.2. This tells us whether R crosses e, and we have a pair of points in the embedding of f representing R, which is exactly what we need to calculate the path of R through the next face.

There are $O(n)$ pairs of a face and an edge of that face, so the total amount of computation before the face sequence repeats periodically is $O(nb)$. (If R ends at a polyhedron vertex before then, we calculate so because R exits a face by two edges at the same time, and we can compare the $O(b)$-bit leaving times in $O(b)$ time.)

Third, we calculate the periodic part of the path. Consider the shape formed by the faces $f_1, f_2, \ldots, f_k$ of $F(R)$ that repeat periodically, as in the bolded part of Figure 6. Copies of this shape attach to each other on copies of a repeated edge e; that is, the entire shape is translated and possibly rotated to identify the copies of e. The composition of the isometries that take f_1 to f_2, f_2 to f_3, and so on is an isometry that takes one copy of e to the next. By Section 1.2, we can check whether the slopes of two copies of e are equal (a case with no rotation) or not.

In the case with no rotation, as in Figure 6, all copies of each edge e are translates of each other by a constant amount, and we can describe all copies of e as line segments from $(x_0 + k\Delta x, y_0 + k\Delta y)$ to $(x_1 + k\Delta x, y_1 + k\Delta y)$ for some $x_0, x_1, y_0, y_1, \Delta x, \Delta y$, and all $k \in \mathbb{N}$.

Then, given the equation $(x - x_1)(y_0 - y_1) = (x_0 - x_1)(y - y_1)$ for R, we can calculate the intersection of R with the lines $(x - x_0)\Delta y = (y - y_0)\Delta x$ and $(x - x_1)\Delta y = (y - y_1)\Delta x$ in a constant number of arithmetic operations. One of those intersections is past the first copy of e and one is before it; without loss of generality, suppose that the one past the first copy of e is at $(x_0 + \kappa\Delta x, y_0 + \kappa\Delta y)$ for some $\kappa \in \mathbb{R}^+$. Then the last copy of e that R intersects is the one corresponding to $k = \lfloor \kappa \rfloor$. We can calculate that for each edge in the repeated sequence of faces (reusing the same calculated composition of isometries). The edge minimizing the resulting values of k (with ties broken by the first edge in the sequence of edges of $F(R)$) is the edge by which R leaves the periodic sequence of faces.

If there is rotation, all copies of each edge e are rotations around a consistent center point $C = (x_C, y_C)$. If the first three copies of one endpoint X of e are $X_0 = (x_0, y_0)$, $X_1 = (x_1, y_1)$, and $X_2 = (x_2, y_2)$, then we can calculate the equations of the bisectors of $\overline{X_0 X_1}$ and $\overline{X_1 X_2}$ in $O(1)$ arithmetic operations, so we can calculate their intersection, which is C. All copies of X are of the form $(x_C, y_C) + \sqrt{(x_0 - x_C)^2 + (y_0 - y_C)^2}(\cos(\theta_0 + k\Delta\theta), \sin(\theta_0 + k\Delta\theta))$ for some θ_0 and θ_1. (We calculate trig functions only precisely enough to take a floor: see below.) Then all copies of X are on the circle $(x - x_C)^2 + (y - y_C)^2 = (x_0 - x_C)^2 + (y_0 - y_C)^2$. We can calculate the (two) intersections of R with that circle in $O(1)$ operations. For each intersection $(x_C, y_C) + \sqrt{(x_0 - x_C)^2 + (y_0 - y_C)^2}(\cos(\theta_0 + \kappa\Delta\theta), \sin(\theta_0 + \kappa\Delta\theta))$, we can calculate $k = \lfloor \kappa \rfloor$ by calculating the first few bits of those trig functions (say, by Taylor expansions). Given such a k, the ray R intersects the kth copy of an edge e, then crosses the circle on which all copies of one endpoint of that edge are. However, if that crossing goes from outside to inside the circle, it may happen that R intersects both the kth and $(k + 1)$st copies of e, even though it left the circle in between them. So, check whether R intersects the $(k + 1)$st copy of e; if so, move on to the next-smallest value of k. There are at most $2n$ endpoints, so after $O(n)$ such operations, we find the first edge on which R leaves the periodic pattern of faces. ◀

▶ **Corollary 2.5.** *A geodesic segment $R(X, \varphi, d)$ can be implicitly represented by $O\left(\frac{d}{\ell}\right)$ subpaths, each of which visits a prefix of a periodic sequence of $O(n)$ faces, which can be computed in $O\left(n\frac{d}{\ell}b\right)$ time.*

Proof. Apply Lemma 2.4 to $R = R\left(X, \varphi, \frac{\ell}{2}\right)$ to generate a point X' and direction φ' of the endpoint of R other than X that is at least distance $\frac{\ell}{2}$ from X and traverses the prefix of some periodic sequence of $O(n)$ faces in $O(nb)$ time. Repeatedly appling Lemma 2.4 to $R\left(X', \varphi', \frac{\ell}{2}\right)$, and again at most $2d/\ell$ times proves the claim. ◀

2.3 Full Algorithm

We are now ready to state the algorithm for finding a closed quasigeodesic in quasipolynomial time:

▶ **Theorem 2.6.** *Let P be a convex polyhedron with n vertices all of curvature at least ε, let L be the length of the longest edge, let ℓ be the smallest distance within a face between a vertex and a nonincident edge, let b be the maximum number of bits of an integer in a constant-size radical expression of a real number representing P. Then, in $O\left(\frac{n^2}{\varepsilon^2}\frac{L}{\ell}b\right)$ time, we can find a closed quasigeodesic on P. The closed quasigeodesic can be implicitly represented by $O\left(\frac{n}{\varepsilon}\right)$ vertex-to-vertex paths, where each path is composed of $O\left(\frac{L}{\ell\varepsilon}\right)$ subpaths each of which visits some prefix of a periodic sequence of $O(n)$ faces.*

Proof. For each vertex V of P, divide the total angle at that vertex (that is, the angles at that vertex in the faces that meet at that vertex) into arcs of size between $\varepsilon/4$ and $\varepsilon/2 < \pi$, making $O(1/\varepsilon)$ such arcs at each vertex.

Construct a directed graph G whose nodes are pairs of a vertex V from P and one of its arcs I, giving the graph $O(n/\varepsilon)$ nodes, with an edge from a node $u = (U, I)$ to a node $v = (V, J)$ if there exists a direction in I such that a quasigeodesic ray starting in that direction from the polyhedron vertex U hits the polyhedron vertex V and can continue from every angle in J.

Let $v = (V, I)$ be a node of G, with corresponding vertex V and arc I spanning angles from φ_1 to φ_2. Compute face sequences for $R_1 = R(V, \varphi_1, L/\varepsilon)$ and $R_2 = R(V, \varphi_2, L/\varepsilon)$ and compare their face sequences $F(R_1)$ and $F(R_2)$. By Lemma 2.3, face sequences $F(R_1)$ and $F(R_2)$ differ somewhere, and their first difference determines a polyhedron vertex reachable in the wedge between R_1 and R_2 via a geodesic from V n a direction between angles φ_1 and φ_2, which can be found by scanning the sequencing. Once we reach such a vertex U, a quasigeodesic can exit the vertex anywhere in an angle equal to that vertex's curvature, which is at least ε, so for at least one of the arcs J of size at most $\varepsilon/2$ at that vertex, the quasigeodesic can exit anywhere in that arc, so we have found an outgoing edge from node v to node $u = (U, J)$.

The preceding algorithm computes an outgoing edge from any node in G, so we repeatedly traverse outgoing edges of G until a node of G is repeated. This cycle in G exactly corresponds to a closed quasigeodesic on the polyhedron, by the definition of the graph at the start of Section 2.1.

This algorithm computes $O(n/\varepsilon)$ edges of G (at most one for every graph node) before finding a cycle. To find an edge, the algorithm computes two face sequences $F(R_1)$ and $F(R_2)$, which by Corollary 2.5 can each be implicitly represented by $O\left(\frac{L}{\ell\varepsilon}\right)$ subpaths, each of which visits a prefix of a periodic sequence of $O(n)$ faces and can be computed in $O\left(\frac{n}{\varepsilon}\frac{L}{\ell}b\right)$ time. Then the geodesic corresponding to each edge of G can be computed through the longest common prefix of these face sequences in the same amount of time. Thus the whole geodesic can be described by $O(n/\varepsilon)$ such vertex-to-vertex paths, and can be constructed in $O\left(\frac{n^2}{\varepsilon^2}\frac{L}{\ell}b\right)$ time, as desired.   ◀

If D is the greatest diameter of a face, then a closed quasigeodesic found by Theorem 2.6 has length $O\left(\frac{n}{\varepsilon}\left(\frac{L}{\varepsilon} + D\right)\right)$, because the quasigeodesic visits $O(n/\varepsilon)$ graph nodes and, by Lemma 2.3, goes a distance at most $L/\varepsilon + D$ between each consecutive pair.

3   Conclusion

It has been known for seven decades [19] that every convex polyhedron has a closed quasigeodesic, but our algorithm is the first finite algorithm to find one. We end with some open problems about extending our approach, though they all seem difficult.

▶ **Open Problem 1.** *Theorem 2.6 does not necessarily find a* non-self-intersecting *closed quasigeodesic, even though at least three are guaranteed to exist. Is there an algorithm to find one? In particular, can we find the* shortest *closed quasigeodesic?*

Any approach similar to Theorem 2.6 is unlikely to resolve this, for several reasons:
1. Parts of a quasigeodesic could enter a vertex at infinitely many angles. Theorem 2.6 makes this manageable by grouping similar angles of entry to a vertex, but if similar angles of entry to a vertex are combined, extensions that would be valid for some of them

but invalid for others are treated as invalid for all of them. For instance, a quasigeodesic found by Theorem 2.6 will almost never turn by the maximum allowed at any vertex, since exiting a vertex at the maximum possible turn from one entry angle to the vertex may mean exiting it with more of a turn than allowed for another very close entry angle. So there are some closed quasigeodesics not findable by Theorem 2.6, and those may include non-self-intersecting ones.

2. Given a vertex and a wedge determined by a range of directions from it, we can find *one* vertex in the wedge, but if we wish to find more than one, the problem becomes more complicated. When we seek only one vertex, we only need consider one unfolding of the faces, which the entire wedge stays in until it hits a vertex; when we pass a vertex, the unfoldings on each side of it might be different, so we multiply the size of the problem by 2 every time we pass a vertex. There may, in fact, be exponentially many non-self-intersecting geodesic paths between two vertices: for instance, Aronov and O'Rourke [9] give the example of a doubly covered regular polygon, in which a geodesic path may visit every vertex in order around the cycle but may skip vertices.

▶ **Open Problem 2.** *Theorem 2.6 is polynomial in not just n but the smallest curvature at a vertex, the length of the longest edge, and the shortest distance within a face between a vertex and an edge not containing it. Are all of those necessary? Can the last be simplified to the length of the shortest side?*

▶ **Open Problem 3.** *Can the algorithm of Theorem 2.6 be extended to nonconvex polyhedra?*

▶ **Open Problem 4.** *Is there an algorithm to find a closed quasigeodesic passing through a number of faces bounded by a polynomial function of n, ε, L, ℓ, and perhaps the minimum total angle of a polyhedron vertex? Does Theorem 2.6 already have such a bound?*

A single quasigeodesic ray may pass through a number of faces not bounded by a function of those parameters before ceasing to cycle periodically: for instance, the geodesic ray of Figure 4 does. However, we have no example for which a whole geodesic wedge passes through a number of faces not bounded by a function of those parameters before containing a vertex.

──── **References** ────

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

2 Werner Ballmann, Gudlaugur Thorbergsson, and Wolfgang Ziller. On the existence of short closed geodesics and their stability properties. In *Seminar on Minimal Submanifolds*, pages 53–63. Princeton University Press, 1983. URL: `https://www.researchgate.net/profile/Wolfgang_Ziller/publication/268500145_On_the_existence_of_short_closed_geodesics_and_their_stability_properties/links/58de87a7a6fdcc41bf8e987f/On-the-existence-of-short-closed-geodesics-and-their-stability-properties.pdf`.

3 Amir M. Ben-Amram. What is a "pointer machine"? *SIGACT News*, 26(2):88–95, June 1995. `doi:10.1145/202840.202846`.

4 A. Bertoni, G. Mauri, and N. Sabadini. Simulations among classes of random access machines and equivalence among numbers succinctly represented. *Annals of Discrete Mathematics*, 25:65–90, 1985.

5 George D. Birkhoff. *Dynamical Systems*, volume 9 of *Colloquium Publications*. American Mathematical Society, 1927. URL: `https://bookstore.ams.org/coll-9`.

6 Christoph Burnikel, Stefan Funke, Kurt Mehlhorn, Stefan Schirra, and Susanne Schmitt. A separation bound for real algebraic expressions. In *Proceedings of the 9th Annual European Symposium on Algorithms*, volume 2161 of *Lecture Notes in Computer Science*, pages 254–265, Aarhus, Denmark, August 2001. URL: `http://graphics.stanford.edu/~sfunke/Papers/ESA01/sepbound01.pdf`.

7   Timothy M. Chan and Mihai Pǎtraşcu. Transdichotomous results in computational geometry, I: Point location in sublogarithmic time. *SIAM Journal on Computing*, 39(2):703–729, 2009. `doi:10.1137/07068669X`.

8   Timothy M. Chan and Mihai Pǎtraşcu. Transdichotomous results in computational geometry, II: Offline search, 2010. Originally published at STOC 2007. `arXiv:1010.1948`.

9   Erik D. Demaine and Joseph O'Rourke. Geodesics: Lyusternik-Schnirelmann. In *Geometric Folding Algorithms: Linkages, Origami, Polyhedra*, section 24.4, pages 372–375. Cambridge University Press, Cambridge, 2007.

10 James R. Driscoll, Neil Sarnak, Daniel D. Sleator, and Robert E. Tarjan. Making data structures persistent. *Journal of Computer and System Sciences*, 38(1):86–124, 1989.

11 Michael L. Fredman and Dan E. Willard. Surpassing the information theoretic bound with fusion trees. *Journal of Computer and System Sciences*, 47(3):424–436, 1993. `doi:10.1016/0022-0000(93)90040-4`.

12 David Harvey and Joris van der Hoeven. Integer multiplication in time $O(n \log n)$. HAL Preprint hal-02070778, 2019. URL: `https://hal.archives-ouvertes.fr/hal-02070778`.

13 Giuseppe F. Italiano and Rajeev Raman. Topics in data structures. In Mikhail J. Atallah and Marina Blanton, editors, *Algorithms and Theory of Computation Handbook*, volume 1, chapter 5, pages 5-1–5-29. CRC Press, second edition, 2010.

14 Jin-Ichi Itoh, Joël Rouyer, and Costin Vîlcu. Polyhedra with simple dense geodesics. *Differential Geometry and its Applications*, 66:242–252, 2019. `doi:10.1016/j.difgeo.2019.07.001`.

15 Daniel Kane, Gregory N. Price, and Erik D. Demaine. A pseudopolynomial algorithm for Alexandrov's Theorem. In *Proceedings of the 11th Algorithms and Data Structures Symposium*, volume 5664 of *Lecture Notes in Computer Science*, pages 435–446, Banff, Canada, August 2009.

16 Donald E. Knuth. *The Art of Computer Programming*, volume 2. Addison-Wesley, 1969.

17 Lazar Lyusternik and Lev Schnirelmann. Sur le probléme de trois géodésiques fermées sur les surfaces de genre 0. *Comptes Rendus de l'Académie des Sciences de Paris*, 189:269–271, 1929.

18 Joseph O'Rourke. Personal communication, 2020.

19 Aleksei Vasilevich Pogorelov. Quasi-geodesic lines on a convex surface. *Matematicheskii Sbornik*, 25(62):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 Arnold Schönhage. On the power of random access machines. In *Proceedings of the 6th International Colloquium on Automata, Languages, and Programming*, volume 71 of *Lecture Notes in Computer Science*, pages 520–529, 1979.

22 Michael Ian Shamos. *Computational Geometry*. PhD thesis, Yale University, 1978.

23 Vikram Sharma and Chee K. Yap. Robust geometric computation. In Jacob E. Goodman, Joseph O'Rourke, and Csaba D. Tóth, editors, *Handbook of Discrete and Computational Geometry*, chapter 45, pages 1189–1223. CRC Press, third edition, 2018.