A PSEUDOPOLYNOMIAL ALGORITHM FOR ALEXANDROV'S THEOREM

DANIEL KANE, GREGORY NATHAN PRICE, AND ERIK D. DEMAINE

ABSTRACT. Alexandrov's Theorem states that every metric with the global topology and local geometry required of a convex polyhedron is in fact the intrinsic metric of some convex polyhedron. Recent work by Bobenko and Izmestiev describes a differential equation whose solution is the polyhedron corresponding to a given metric. We describe an algorithm based on this differential equation to compute the polyhedron given the metric, and prove a pseudopolynomial bound on its running time.

Dagstuhl Seminar Proceedings 09111
Computational Geometry
http://drops.dagstuhl.de/opus/volltexte/2009/2032

1. Introduction

Consider the intrinsic metric M induced on the surface of a convex body in $\mathbb{R}^3$. Clearly M is homeomorphic to a sphere, and locally convex in the sense that a circle of radius r has circumference at most $2\pi r$.

In 1949, Alexandrov and Pogorelov [1] proved that these two necessary conditions are actually sufficient: every metric M that is homeomorphic to a 2-sphere and locally convex can be embedded as the surface of a convex body in $\mathbb{R}^3$. Because Alexandrov and Pogorelov's proof is highly nonconstructive, their work opened the question of how to produce the embedding given a concrete M.

To enable computation we require that M be a polyhedral metric, locally isometric to $\mathbb{R}^2$ at all but n points (vertices). Now the theorem is that every polyhedral metric, a complex of triangles with the topology of a sphere and positive curvature at each vertex, can be embedded as an actual convex polyhedron in $\mathbb{R}^3$. This case of the Alexandrov-Pogorelov theorem was proven by Alexandrov in 1941 [1], also nonconstructively. Further, Cauchy showed in 1813 [3] that such an embedding must be unique. All the essential geometry of the general case is preserved in the polyhedral case, because every metric satisfying the general hypothesis can be polyhedrally approximated.

In 1996, Sabitov [11, 10, 12, 5] showed how to enumerate all the isometric maps $M \to \mathbb{R}^3$ for a polyhedral metric M, so that one could carry out this enumeration and identify the one map that gives a convex polyhedron. In 2005, Fedorchuk and Pak [6] showed an exponential upper bound on the number of such maps. An exponential lower bound is easy to find, so this algorithm takes time exponential in n and is therefore unsatisfactory.

Lubiw and O'Rourke [8] showed in 1996 how to solve efficiently the complementary problem of deciding whether a plane polygon can be glued along its edges to produce a valid M. In 2002, Demaine, Demaine, Lubiw, and O'Rourke [4] showed that some polygons can be glued to themselves along their boundaries to form exponentially many combinatorially distinct M and uncountably many nonisometric M. Both papers left open the problem of constructing the polyhedron that results from any such gluing.

The textbook by Demaine and O'Rourke [5] describes in detail each step in this history, and poses again the problem of constructing the embedding given M.

Recent work by Bobenko and Izmestiev [2] produced a new proof of Alexandrov's Theorem, describing a certain ordinary differential equation (ODE) and initial conditions whose solution contains sufficient information to construct the embedding by elementary geometry. This work included a computer implementation of the ODE, which empirically produces accurate approximations of embeddings of metrics on which it is tested.

In this work, we describe an algorithm based on the Bobenko-Izmestiev ODE, and prove a pseudopolynomial bound on its running time. Specifically, call an embedding of M *ε-accurate* if the metric is distorted by at most $1+\varepsilon$, and *ε-convex* if each dihedral angle is at most $\pi + \varepsilon$. Then we show the following theorem:

Theorem 1.1. *Given a polyhedral metric M with n vertices, ratio S between the diameter and the smallest distance between vertices, and defect (discrete Gaussian curvature) between ε_1 and $2\pi - \varepsilon_8$ at each vertex, an ε_6-accurate ε_9-convex* *embedding of M can be found in time $O\left(n^{913/2}S^{831}/(\varepsilon^{121}\varepsilon_1^{445}\varepsilon_8^{616})\right)$ where $\varepsilon = \min(\varepsilon_6/nS, \varepsilon_9\varepsilon_1^2/nS^6)$.*

The exponents in the time bound of Theorem 1.1 are remarkably large. Thankfully, no evidence suggests our algorithm actually takes as long to run as the bound allows. On the contrary, our analysis relies on bounding approximately a dozen geometric quantities, and to keep the analysis tractable we use the simplest bound whenever available. The algorithm's actual performance is governed by the actual values of these quantities, and therefore by whatever sharper bounds can be proven by a stingier analysis.

To describe our approach, consider an embedding of the metric M as a convex polyhedron in $\mathbb{R}^3$, and choose an arbitrary origin O in the surface's interior. Then it is not hard to see that the n distances $r_i = \overline{Ov_i}$ from the origin to the vertices v_i, together with M and the combinatorial data describing which polygons on M are faces of the polyhedron, suffice to reconstruct the embedding: the tetrahedron formed by O and each triangle is rigid in $\mathbb{R}^3$, and we have no choice in how to glue them to each other. In Lemmas 3.2 and 3.3 below, we show that in fact the radii alone suffice to reconstruct the embedding, to do so efficiently, and to do so even with radii of finite precision.

Therefore in order to compute the unique embedding of M that Alexandrov's Theorem guarantees exists, we compute a set of radii $r = \{r_i\}_i$ and derive a triangulation T. The exact radii satisfy three conditions:

(1) the radii r give valid tetrahedra from O to each face of T;
(2) with these tetrahedra, the dihedral angles at each exterior edge total at most π; and
(3) with these tetrahedra, the dihedral angles about each radius sum to 2π.

In our computation, we begin with a set of large initial radii $r_i = R$ satisfying Conditions 1 and 2, and write $\kappa = \{\kappa_i\}_i$ for the differences by which Condition 3 fails about each radius. We then iteratively adjust the radii to bring κ near zero and satisfy Condition 3 approximately, maintaining Conditions 1 and 2 throughout.

The computation takes the following form. We describe the Jacobian $\left(\frac{\partial\kappa_i}{\partial r_j}\right)_{ij}$, showing that it can be efficiently computed and that its inverse is pseudopolynomially bounded. We show further that the Hessian $\left(\frac{\partial\kappa_i}{\partial r_j \partial r_k}\right)_{ijk}$ is also pseudopolynomially bounded. It follows that a change in r in the direction of smaller κ as described by the Jacobian, with some step size only pseudopolynomially small, makes progress in reducing $|\kappa|$. The step size can be chosen online by doubling and halving, so it follows that we can take steps of the appropriate size, pseudopolynomial in number, and obtain an r that zeroes κ to the desired precision in pseudopolynomial total time. Theorem 1.1 follows.

The construction of [2] is an ODE in the same n variables r_i, with a similar starting point and with the derivative of r driven similarly by a desired path for κ. Their proof differs in that it need only show existence, not a bound, for the Jacobian's inverse, in order to invoke the inverse function theorem. Similarly, while we must show a pseudopolynomial lower bound (Lemma 6.7) on the altitudes of the tetrahedra during our computation, the prior work shows only that these altitudes remain positive. In general our computation requires that the known open conditions—this quantity is positive, that map is nondegenerate—be replaced by stronger compact conditions—this quantity is lower-bounded, that map's inverse is bounded. We model our proofs of these strengthenings on the proofs in [2] of the simpler open conditions, and we directly employ several other results from that paper where possible.

The remainder of this paper supplies the details of the proof of Theorem 1.1. We give background in Section 2, and detail the main argument in Section 3. We bound the Jacobian in Section 4 and the Hessian in Section 5. Finally, some lemmas are deferred to Section 6 for clarity.

2. Background and Notation

In this section we define our major geometric objects and give the basic facts about them. We also define some parameters describing our central object that we will need to keep bounded throughout the computation.

2.1. **Geometric notions.** Central to our argument are two dual classes of geometric structures introduced by Bobenko and Izmestiev in [2] under the names of "generalized convex polytope" and "generalized convex polyhedron". Because in other usages the distinction between "polyhedron" and "polytope" is that a polyhedron is a three-dimensional polytope, and because both of these objects are three-dimensional, we will refer to these objects as "generalized convex polyhedra" and "generalized convex dual polyhedra" respectively to avoid confusion.

First, we define the objects that our main theorem is about.

Definition 2.1. A metric M homeomorphic to the sphere is a *polyhedral metric* if each $x \in M$ has an open neighborhood isometric either to a subset of $\mathbb{R}^2$ or to a cone of angle less than 2π with x mapped to the apex, and if only finitely many x, called the *vertices* $V(M) = \{v_i\}_i$ of M, fall into the latter case.

The *defect* δ_i at a vertex $v_i \in V(M)$ is the difference between 2π and the total angle at the vertex, which is positive by the definition of a vertex.

An *embedding* of M is a continuous map $f : M \to \mathbb{R}^3$. An embedding f is *ε-accurate* if it distorts the metric M by at most $1 + \varepsilon$, and *ε-convex* if $f(M)$ is a polyhedron and each dihedral angle in $f(M)$ is at most $\pi + \varepsilon$.

A *perfect embedding* of a polyhedral metric M is an isometry $f : M \to \mathbb{R}^3$ such that $f(M)$ is a convex polyhedron. Equivalently, an embedding is perfect if 0-accurate and 0-convex.

Alexandrov's Theorem is that every polyhedral metric has a unique perfect embedding, and our contribution is a pseudopolynomial-time algorithm to construct ε-accurate ε-convex embeddings as approximations to this perfect embedding.

Definition 2.2. In a tetrahedron $ABCD$, write $\angle CABD$ for the dihedral angle along edge AB.

Definition 2.3. A *triangulation* of a polyhedral metric M is a decomposition into Euclidean triangles whose vertex set is $V(M)$. Its vertices are denoted by $V(T) = V(M)$, its edges by $E(T)$, and its faces by $F(T)$.

A *radius assignment* on a polyhedral metric M is a map $r : V(M) \to \mathbb{R}_+$. For brevity we write r_i for $r(v_i)$.

Given a polyhedral metric M, a triangulation T, and a radius assignment r, the *generalized convex polyhedron* $P = (M, T, r)$ is a certain metric space on the topological cone on M with an apex O, if M, T, r are suitable. Let the cone OF for each face $F \in F(T)$ be isometric to a Euclidean tetrahedron with base F and side edges given by r. We require that the total dihedral angle about each edge of T be at most π, and about each edge Ov_i at most 2π.

Write $\kappa_i \triangleq 2\pi - \sum_{jk} \angle v_j O v_i v_k$ for the curvature about Ov_i, and $\phi_{ij} \triangleq \angle v_i O v_j$ for the angle between vertices v_i, v_j seen from the apex.

Our algorithm, following the construction in [2], will choose a radius assignment for the M in question and iteratively adjust it until the associated generalized convex polyhedron P fits nearly isometrically in $\mathbb{R}^3$. The resulting radii will give an ε-accurate ε-convex embedding of M into $\mathbb{R}^3$.

In the argument we will require several geometric objects related to generalized convex polyhedra.

Definition 2.4. A *Euclidean simplicial complex* is a metric space on a simplicial complex where the metric restricted to each cell is Euclidean.

A *generalized convex polygon* is a Euclidean simplicial 2-complex homeomorphic to a disk, where all triangles have a common vertex V, the total angle at V is no more than 2π, and the total angle at each other vertex is no more than π.

Given a generalized convex polyhedron $P = (M, T, r)$, the corresponding *generalized convex dual polyhedron* $D(P)$ is a certain Euclidean simplicial 3-complex. Let O be an vertex called the *apex*, A_i a vertex with $OA_i = h_i \triangleq 1/r_i$ for each i.

For each edge $v_iv_j \in E(T)$ bounding triangles $v_iv_jv_k$ and $v_jv_iv_l$, construct two simplices $OA_iA_{jil}A_{ijk}$, $OA_jA_{ijk}A_{jil}$ in $D(P)$ as follows. Embed the two tetrahedra $Ov_iv_jv_k, Ov_jv_iv_l$ in $\mathbb{R}^3$. For each $i' \in \{i, j, k, l\}$, place $A_{i'}$ along ray $Ov_{i'}$ at distance $h_{i'}$, and draw a perpendicular plane $P_{i'}$ through the ray at $A_{i'}$. Let A_{ijk}, A_{jil} be the intersection of the planes P_i, P_j, P_k and P_j, P_i, P_l respectively.

Now identify the vertices $A_{ijk}, A_{jki}, A_{kij}$ for each triangle $v_iv_jv_k \in F(T)$ to produce the Euclidean simplicial 3-complex $D(P)$. Since the six simplices produced about each of these vertices A_{ijk} are all defined by the same three planes P_i, P_j, P_k with the same relative configuration in $\mathbb{R}^3$, the total dihedral angle about each OA_{ijk} is 2π. On the other hand, the total dihedral angle about OA_i is $2\pi - \kappa_i$, and the face about A_i is a generalized convex polygon of defect κ_i. Let

$$h_{ij} = \frac{h_j - h_i \cos\phi_{ij}}{\sin\phi_{ij}}$$

be the altitude in this face from its apex A_i to side $A_{ijk}A_{jil}$.

Definition 2.5. A *singular spherical polygon* (or *triangle*, *quadrilateral*, etc) is a simplexwise spherical metric space on a 2-complex homeomorphic to a disk, where the total angle at each interior vertex is at most 2π. A singular spherical polygon is *convex* if the total angle at each boundary vertex is at most π.

A *singular spherical metric* is a simplexwise spherical metric space on a 2-complex homeomorphic to a sphere, where the total angle at each vertex is at most 2π.

The Jacobian bound in Section 4 makes use of certain multilinear forms described in [2].

Definition 2.6. The *dual volume* $\mathrm{vol}(h)$ is the volume of the generalized convex dual polyhedron $D(P)$, a cubic form in the dual altitudes h.

The *mixed volume* $\mathrm{vol}(\cdot,\cdot,\cdot)$ is the symmetric trilinear form that formally extends the cubic form $\mathrm{vol}(\cdot)$:

$$\mathrm{vol}(a,b,c) \triangleq \frac{1}{6}(\mathrm{vol}(a+b+c)-\mathrm{vol}(a+b)-\mathrm{vol}(b+c)-\mathrm{vol}(c+a)+\mathrm{vol}(a)+\mathrm{vol}(b)+\mathrm{vol}(c)).$$

The ith *dual face area* $E_i(g(i))$ is the area of the face around A_i in $D(P)$, a quadratic form in the altitudes $g(i) \triangleq \{h_{ij}\}_j$ within this face.

The ith *mixed area* $E_i(\cdot,\cdot)$ is the symmetric bilinear form that formally extends the quadratic form $E_i(\cdot)$:

$$E_i(a,b) \triangleq \frac{1}{2}(E_i(a+b) - E_i(a) - E_i(b)).$$

Let π_i be the linear map

$$\pi_i(h)_j \triangleq \frac{h_j - h_i \cos\phi_{ij}}{\sin\phi_{ij}}$$

so that $\pi_i(h) = g(i)$. Then define

$$F_i(a,b) \triangleq E_i(\pi_i(a), \pi_i(b)).$$

so that $F_i(h,h) = E_i(g(i), g(i))$ is the area of face i.

Observe that $\mathrm{vol}(h,h,h) = \frac{1}{3}\sum_i h_i F_i(h,h)$, so that by a simple computation

$$\mathrm{vol}(a,b,c) = \frac{1}{3}\sum_i a_i F_i(b,c).$$

2.2. **Weighted Delaunay triangulations.** The triangulations we require at each step of the computation are the weighted Delaunay triangulations used in the construction of [2]. We give a simpler definition inspired by Definition 14 of [7].

Definition 2.7. The *power* $\pi_v(p)$ of a point p against a vertex v in a polyhedral metric M with a radius assigment r is $pv^2 - r(v)^2$.

The *center* $C(v_iv_jv_k)$ of a triangle $v_iv_jv_k \in T(M)$ when embedded in $\mathbb{R}^2$ is the unique point p such that $\pi_{v_i}(p) = \pi_{v_j}(p) = \pi_{v_k}(p)$, which exists by the radical axis theorem from classical geometry. The quantity $\pi_{v_i}(p) = \pi(v_iv_jv_k)$ is the *power* of the triangle.

A triangulation T of a polyhedral metric M with radius assignment r is *locally convex* at edge v_iv_j with neighboring triangles $v_iv_jv_k, v_jv_iv_l$ if $\pi_{v_l}(C(v_iv_jv_k)) > \pi_{v_l}(v_k)$ and $\pi_{v_k}(C(v_jv_iv_l)) > \pi_{v_k}(v_l)$ when $v_iv_jv_k, v_jv_iv_l$ are embedded together in $\mathbb{R}^2$.

A *weighted Delaunay triangulation* for a radius assignment r on a polyhedral metric M is a triangulation T that is locally convex at every edge.

A weighted Delaunay triangulation can be computed in time $O(n^2 \log n)$ by a simple modification of the "continuous Dijkstra" algorithm of [9].

The radius assignment r and triangulation T admits a tetrahedron $Ov_iv_jv_k$ just if the power of $v_iv_jv_k$ is negative, and the squared altitude of O in this tetrahedron is $-\pi(v_iv_jv_k)$. The edge v_iv_j is convex when the two neighboring tetrahedra are embedded in $\mathbb{R}^3$ just if it is locally convex in the triangulation as in Definition 2.7. A weighted Delaunay triangulation with negative powers therefore gives a valid generalized convex polyhedron if the curvatures κ_i are positive. For each new radius assignment r in the computation of Section 3 we therefore compute the weighted Delaunay triangulation and proceed with the resulting generalized convex polyhedron, in which Lemma 6.7 guarantees a positive altitude and the choices in the computation guarantee positive curvatures.

2.3. **Notation for bounds.**

Definition 2.8. Let the following bounds be observed:

(1) n is the number of vertices on M. By Euler's formula, $|E(T)|$ and $|F(T)|$ are both $O(n)$.
(2) $\varepsilon_1 \triangleq \min_i \delta_i$ is the minimum defect.
(3) $\varepsilon_2 \triangleq \min_i(\delta_i - \kappa_i)$ is the minimum defect-curvature gap.
(4) $\varepsilon_3 \triangleq \min_{ij \in E(T)} \phi_{ij}$ is the minimum angle between radii.
(5) $\varepsilon_4 \triangleq \max_i \kappa_i$ is the maximum curvature.
(6) $\varepsilon_5 \triangleq \min_{v_i v_j v_k \in F(T)} \angle v_i v_j v_k$ is the smallest angle in the triangulation. Observe that obtuse angles are also bounded: $\angle v_i v_j v_k < \pi - \angle v_j v_i v_k \leq \pi - \varepsilon_5$.
(7) ε_6 is used for the desired accuracy in embedding M.
(8) $\varepsilon_7 \triangleq (\max_i \frac{\kappa_i}{\delta_i})/(\min_i \frac{\kappa_i}{\delta_i}) - 1$ is the extent to which the ratio among the κ_i varies from that among the δ_i. We will keep $\varepsilon_7 < \varepsilon_8/4\pi$ throughout the computation.
(9) $\varepsilon_8 \triangleq \min_i(2\pi - \delta_i)$ is the minimum angle around a vertex, the complement of the maximum defect.
(10) ε_9 is used for the desired approximation to convexity in embedding M.
(11) D is the diameter of M.
(12) ℓ is the shortest distance $v_i v_j$ between vertices.
(13) $S \triangleq D/\ell$ is the maximum ratio of distances.
(14) $d_0 \triangleq \min_{p \in M} Op$ is the minimum height of the apex off of any point on M.
(15) $d_1 \triangleq \min_{v_i v_j \in E(T)} d(O, v_i v_j)$ is the minimum distance from the apex to any edge of T.
(16) $d_2 \triangleq \min_i r_i$ is the minimum distance from the apex to any vertex of M.
(17) $H \triangleq 1/d_0$; the name is justified by $h_i = 1/r_i \leq 1/d_0$.
(18) $R \triangleq \max_i r_i$, so $1/H \leq r_i \leq R$ for all i.
(19) $T \triangleq HR$ is the maximum ratio of radii.

Of these bounds, $n, \varepsilon_1, \varepsilon_8, S$ are fundamental to the given metric M, and D a dimensionful parameter given by M. The values $\varepsilon_6, \varepsilon_9$ define the objective to be achieved, and our computation will drive ε_4 toward zero while maintaining ε_2 large and ε_7 small. In Section 6 we bound the remaining parameters $\varepsilon_3, \varepsilon_5, R, d_0, d_1, d_2$ in terms of these.

Definition 2.9. Let $\mathbf{J}$ denote the Jacobian $\left(\frac{\partial \kappa_i}{\partial r_j}\right)_{ij}$, and $\mathbf{H}$ the Hessian $\left(\frac{\partial \kappa_i}{\partial r_j \partial r_k}\right)_{ijk}$.

3. Main Theorem

In this section, we prove our main theorem using the results proved in the remaining sections. Recall

Theorem 1.1. *Given a polyhedral metric M with n vertices, ratio S (the* spread*) between the diameter and the smallest distance between vertices, and defect at least*

ε_1 and at most $2\pi-\varepsilon_8$ at each vertex, an ε_6-accurate ε_9-convex embedding of M can be found in time $O\left(n^{913/2}S^{831}/(\varepsilon^{121}\varepsilon_1^{445}\varepsilon_8^{616})\right)$ where $\varepsilon = \min(\varepsilon_6/nS, \varepsilon_9\varepsilon_1^2/nS^6)$.

The algorithm of Theorem 1.1 obtains an approximate embedding of the polyhedral metric M in $\mathbb{R}^3$. Its main subroutine is described by the following theorem:

Theorem 3.1. *Given a polyhedral metric M with n vertices, ratio S (the* spread*) between the diameter and the smallest distance between vertices, and defect at least ε_1 and at most $2\pi-\varepsilon_8$ at each vertex, a radius assignment r for M with maximum curvature at most ε can be found in time $O\left(n^{913/2}S^{831}/(\varepsilon^{121}\varepsilon_1^{445}\varepsilon_8^{616})\right)$.*

Proof. Let a *good* assignment be a radius assignment r that satisfies two bounds: $\varepsilon_7 < \varepsilon_8/4\pi$ so that Lemmas 6.5–6.7 apply and r therefore by the discussion in Subsection 2.2 produces a valid generalized convex polyhedron for M, and $\varepsilon_2 = \Omega(\varepsilon_1^2\varepsilon_8^3/n^2S^2)$ on which our other bounds rely. By Lemma 6.1, there exists a good assignment r^0. We will iteratively adjust r^0 through a sequence r^t of good assignments to arrive at an assignment r^N with maximum curvature $\varepsilon_4^N < \varepsilon$ as required. At each step we recompute T as a weighted Delaunay triangulation according to Subsection 2.2.

Given a good assignment $r = r^n$, we will compute another good assignment $r' = r^{n+1}$ with $\varepsilon_4 - \varepsilon_4' = \Omega\left(\varepsilon_1^{445}\varepsilon_4^{121}\varepsilon_8^{616}/(n^{907/2}S^{831})\right)$. It follows that from r^0 we can arrive at a satisfactory r^N with $N = O\left((n^{907/2}S^{831})/(\varepsilon^{121}\varepsilon_1^{445}\varepsilon_8^{616})\right)$.

To do this, let $\mathbf{J}$ be the Jacobian $(\frac{\partial\kappa_i}{\partial r_j})_{ij}$ and $\mathbf{H}$ the Hessian $\left(\frac{\partial\kappa_i}{\partial r_j\partial r_k}\right)_{ijk}$, evaluated at r. The goodness conditions and the objective are all in terms of κ, so we choose a desired new curvature vector κ^* in κ-space and apply the inverse Jacobian to get a new radius assignment $r' = r + \mathbf{J}^{-1}(\kappa^* - \kappa)$ in r-space. The actual new curvature vector κ' differs from κ^* by an error at most $\frac{1}{2}|\mathbf{H}||r'-r|^2 \le \left(\frac{1}{2}|\mathbf{H}||\mathbf{J}^{-1}|^2\right)|\kappa^*-\kappa|^2$, quadratic in the desired change in curvatures with a coefficient

$$C \triangleq \frac{1}{2}|\mathbf{H}||\mathbf{J}^{-1}|^2 = O\left(\frac{n^{3/2}S^{14}}{\varepsilon_5^3}\frac{R^{23}}{D^{14}d_0^3d_1^8}\left(\frac{n^{7/2}T^2}{\varepsilon_2\varepsilon_3^3\varepsilon_4}R\right)^2\right) = O\left(\frac{n^{905/2}S^{831}}{\varepsilon_1^{443}\varepsilon_4^{121}\varepsilon_8^{616}}\right)$$

by Theorems 4.1 and 5.2 and Lemmas 6.3, 6.1, 6.7, and 6.4.

Therefore pick a step size p, and choose κ^* such that

$$\kappa_i^* - \kappa_i = -p\kappa_i - p\left(\kappa_i - \delta_i \min_j \frac{\kappa_j}{\delta_j}\right). \tag{1}$$

Consider a hypothetical r^* that gives the curvatures κ^*, and examine the conditions on $\varepsilon_4, \varepsilon_2, \varepsilon_7$ in turn.

Both terms on the right-hand side of (1) are nonpositive, so each κ_i decreases by at least $p\kappa_i$. Therefore the maximum curvature ε_4 decreases by at least $p\varepsilon_4$. If any defect-curvature gap $\delta_i - \kappa_i$ is less than $\varepsilon_1/2$, then it increases by at least $p\kappa_i \ge p(\delta_i - \varepsilon_1/2) \ge p(\varepsilon_1/2)$; so the minimum defect-curvature gap ε_2 either increases by at least $p\varepsilon_1/2$ or is at least $\varepsilon_1/2$ already. Finally, the $-p\kappa_i$ term decreases each κ_i in the same ratio and therefore preserves ε_7, and the $-p\left(\kappa_i - \delta_i \min_j(\kappa_j/\delta_j)\right)$ term decreases each ratio κ_i/δ_i by p times the difference $(\kappa_i/\delta_i - \min_j(\kappa_j/\delta_j))$ and therefore reduces ε_7 by $p\varepsilon_7$. Therefore κ^* would satisfy all three conditions with some room to spare.

In particular, if we choose p to guarantee that each κ'_i differs from κ^*_i by at most $p\varepsilon_4/2$, at most $p\varepsilon_1/2$, and at most $p(\varepsilon_1/4\pi)(\min_i \kappa_i)$, then this discussion shows that the step from r to r' will make at least half the ideal progress $p\varepsilon_4$ in ε_4 and keep $\varepsilon_2, \varepsilon_7$ within bounds.

Since

$$\min_i \kappa_i \geq (\max_j \kappa_j)(\min_{ij} \delta_i/\delta_j)(1+\varepsilon_7)^{-1} \geq \varepsilon_4(\varepsilon_1/2\pi)/2 = \varepsilon_1\varepsilon_4/4\pi$$

and since

$$|\kappa' - \kappa^*|_\infty \leq |\kappa' - \kappa^*| \leq C|\kappa^* - \kappa|^2 \leq 4Cp^2|\kappa|^2 \leq 4Cp^2 n\varepsilon_4^2$$

this can be done by choosing

$$p = \varepsilon_1^2/64\pi^2 n\varepsilon_4 C, \tag{2}$$

which produces a good radius assignment r' in which ε_4 has declined by at least

$$\frac{p\varepsilon_4}{2} = \frac{\varepsilon_1^2}{128\pi^2 nC} = \Omega\left(\frac{\varepsilon_1^{445}\varepsilon_4^{121}\varepsilon_8^{616}}{n^{907/2}S^{831}}\right)$$

as required. Any smaller p will also produce a good assignment r' and decrease ε_4 by at least $p\varepsilon_4/2$ proportionally.

As a simplification, we need not compute p exactly according to (2). Rather, we choose the step size p^t at each step, trying first p^{t-1} (with p^0 an arbitrary constant) and computing the actual curvature error $|\kappa'-\kappa^*|$. If the error exceeds its maximum acceptable value $p\varepsilon_1^2\varepsilon_4/16\pi^2$ then we halve p^t and try step t again, and if it falls below half this value then we double p^t for the next round. Since we double at most once per step and halve at most once per doubling plus a logarithmic number of times to reach an acceptable p, this doubling and halving costs only a constant factor. Even more important than the resulting simplification of the algorithm, this technique holds out the hope of actual performance exceeding the proven bounds.

Now each of the N iterations of the computation go as follows. Compute the weighted Delaunay triangulation T^t for r^t in time $O(n^2 \log n)$ as described in Subsection 2.2. Compute the Jacobian $\mathbf{J}^t$ in time $O(n^2)$ using formulas (14, 15) in [2]. Choose a step size p^t, possibly adjusting it, as discussed above. Finally, take the resulting r' as r^{t+1} and continue. The computation of κ^* to check p^t runs in linear time, and that of r' in time $O(n^\omega)$ where $\omega < 3$ is the time exponent of matrix multiplication. Each iteration therefore costs time $O(n^3)$, and the whole computation costs time $O(n^3 N)$ as claimed. □

Now with our radius assignment r for M and the resulting generalized convex polyhedron P with curvatures all near zero, it remains to approximately embed P and therefore M in $\mathbb{R}^3$. As a motivating warmup, we observe that this is easy to do given exact values for r and a model with exact computation:

Lemma 3.2. *In the real computation model with square roots (aka the straightedge-compass model), there is an algorithm that, given a polyhedral metric M of n vertices and a radius assignment r on M that corresponds to a generalized convex polyhedron P with all curvatures zero, produces explicitly by vertex coordinates a perfect embedding of M in time $O(n^2 \log n)$.*

Proof. Compute the weighted Delaunay triangulation for r in time $O(n^2 \log n)$ as described in Subsection 2.2, obtaining the combinatorial structure of P and the side lengths of its tetrahedra. Now each of these tetrahedra is rigid, so embed one in space arbitrarily and embed each neighboring tetrahedron in turn, spending total time $O(n)$. Because the curvatures are zero the tetrahedra embed exactly without gaps. □

In a realistic model, we compute only with bounded precision, and in any case Theorem 3.1 gives us only curvatures near zero, not equal to zero. Lemma 3.3 produces an embedding in this case, settling for less than exact isometry and exact convexity.

Lemma 3.3. *There is an algorithm that, given a radius assignment r for which the corresponding curvatures κ_i are all less than $\varepsilon = O\left(\min(\varepsilon_6/nS, \varepsilon_9\varepsilon_1^2/nS^6)\right)$ for some constant factor, produces explicitly by vertex coordinates in time $O(n^2 \log n)$ an ε_6-accurate ε_9-convex embedding of M.*

Proof. Compute the weighted Delaunay triangulation T of r on M, and consider the tetrahedra OF_i for $F \in F(T)$. Embed some OF_i in $\mathbb{R}^3$ arbitrarily, embed its neighbors next to it, and so forth, leaving gaps as required by the positive curvature, and call this configuration Q. Since the curvature around each radius is less than ε, the several copies of each vertex will be separated by at most $n\varepsilon D$. Now replace the several copies of each vertex by their centroid, so that the tetrahedra are distorted but leave no gaps. Call the resulting polyhedron P and its surface metric M'. The computation of the weighted Delaunay triangulation takes time $O(n^2 \log n)$ as discussed in Subsection 2.2, and the remaining steps require time $O(n)$. We claim this embedding is ε_6-accurate and ε_9-convex.

To show ε_6-accuracy, observe that since each copy of each vertex was moved by at most $n\varepsilon D$ from Q to P, no edge of any triangle was stretched by more than a ratio $n\varepsilon S$, and the piecewise linear map between faces relates M' to M with distortion $n\varepsilon S \le \varepsilon_6$ as required.

Now we show ε_9-convexity. Consider two neighboring triangles $v_iv_jv_k, v_jv_iv_l$ in T; we will show the exterior dihedral angle is at least $-\varepsilon_9$. First, consider repeating the embedding with $Ov_iv_jv_k$ the original tetrahedron, so that $Ov_iv_jv_k, Ov_jv_iv_l$ embed without gaps. This moves each vertex by at most $n\varepsilon D$, and makes the angle $v_lv_iv_jv_k$ convex and the tetrahedron $v_lv_iv_jv_k$ have positive signed volume. The volume of this tetrahedron in the P configuration is therefore at least $-n\varepsilon D^3$, since the derivative of the volume in any vertex is the area of the opposite face, which is at always less than D^2 since the sides remain $(1 + o(1))D$ in length.

Therefore suppose the exterior angle $\angle v_l v_i v_j v_k$ is negative. Then by Lemma 5.4 and Lemma 6.4,

$$\sin \angle v_l v_i v_j v_k = \frac{3}{2} \frac{[v_l v_i v_j v_k][v_i v_j]}{[v_i v_j v_l][v_j v_i v_k]} \geq -\frac{(n\varepsilon D^3)D}{(\ell^2 \varepsilon_5/4)^2} \geq -\varepsilon \frac{576 n S^6}{\varepsilon_2^2}$$

and since $\varepsilon_2 \geq \varepsilon_1/2$ at the end of the computation, $\angle v_l v_i v_j v_k \geq -\varepsilon 2304 n S^6/\varepsilon_1^2 \geq -\varepsilon_9$ as claimed. □

We now have all the pieces to prove our main theorem.

Proof of Theorem 1.1. Let $\varepsilon \triangleq O\left(\min(\varepsilon_6/nS, \varepsilon_9 \varepsilon_1^2/nS^6)\right)$, and apply the algorithm of Theorem 3.1 to obtain in time $O\left(n^{913/2} S^{831}/(\varepsilon^{121}\varepsilon_1^{445}\varepsilon_8^{616})\right)$ a radius assignment r for M with maximum curvature $\varepsilon_4 \leq \varepsilon$.

Now apply the algorithm of Lemma 3.3 to obtain in time $O(n^2 \log n)$ the desired embedding and complete the computation. □

4. Bounding the Jacobian

Theorem 4.1. *The Jacobian* $\mathbf{J}$*'s inverse is pseudopolynomially bounded by* $|\mathbf{J}^{-1}| = O\left(\frac{n^{7/2}T^2}{\varepsilon_2\varepsilon_3^3\varepsilon_4} R\right)$.

Proof. Our argument parallels that of Corollary 2 in [2], which concludes that the same Jacobian is nondegenerate. Theorem 4 of [2] shows that this Jacobian equals the Hessian of the volume of the dual $D(P)$. The meat of the corollary's proof is in Theorem 5 of [2], which begins by equating this Hessian to the bilinear form $6\,\mathrm{vol}(h, \cdot, \cdot)$ derived from the mixed volume we defined in Definition 2.6. So we have to bound the inverse of this bilinear form.

To do this it suffices to show that the form $\mathrm{vol}(h, x, \cdot)$ has norm at least $\Omega\left(\frac{\varepsilon_2\varepsilon_3^3\varepsilon_4}{n^{7/2}T^2} \frac{|x|}{R}\right)$ for all vectors x. Equivalently, suppose some x has $|\mathrm{vol}(h, x, z)| \leq |z|$ for all z; we show $|x| = O\left(\frac{n^{7/2}T^2}{\varepsilon_2\varepsilon_3^3\varepsilon_4} R\right)$.

To do this we follow the proof in Theorem 5 of [2] that the same form $\mathrm{vol}(h, x, \cdot)$ is nonzero for x nonzero. Throughout the argument we work in terms of the dual $D(P)$.

Recall that for each i, $\pi_i x$ is defined as the vector $\{x_{ij}\}_j$. It suffices to show that for all i

$$|\pi_i x|_2^2 = O\left(\frac{n^3 T^3}{\varepsilon_2^2 \varepsilon_3 \varepsilon_4} R^2 + \frac{n^2 T^2}{\varepsilon_2 \varepsilon_3 \varepsilon_4} R|x|_1\right)$$

since then by Lemma 4.2

$$|x|_2^2 \leq \frac{4n}{\varepsilon_3^2} \max_i |\pi_i x|_2^2 = O\left(\frac{n^4 T^3}{\varepsilon_2^2 \varepsilon_3^3 \varepsilon_4} R^2 + \frac{n^3 T^2}{\varepsilon_2 \varepsilon_3^3 \varepsilon_4} R|x|_1\right),$$

and since $|x|_1 \le \sqrt{n}|x|_2$ and $X^2 \le a+bX$ implies $X \le \sqrt{a}+b$, $|x|_2 = O\left(\frac{n^{7/2}T^2}{\varepsilon_2\varepsilon_3^3\varepsilon_4}R\right)$. Therefore fix an arbitrary i, let $g = \pi_i h$ and $y = \pi_i x$, and we proceed to bound $|y|_2$.

We break the space on which E_i acts into the 1-dimensional positive eigenspace of E_i and its $(k-1)$-dimensional negative eigenspace, since by Lemma 3.4 of [2] the signature of E_i is $(1, k-1)$, where k is the number of neighbors of v_i. Write λ_+ for the positive eigenvalue and $-E_i^-$ for the restriction to the negative eigenspace so that E_i^- is positive definite, and decompose $g = g_+ + g_-$, $y = y_+ + y_-$ by projection into these subspaces. Then we have

$$\begin{aligned} G &\triangleq E_i(g,g) = \lambda_+ g_+^2 - E_i^-(g_-, g_-) \triangleq \lambda_+ g_+^2 - G_- \\ E_i(g,y) &= \lambda_+ g_+ y_+ - E_i^-(g_-, y_-) \\ Y &\triangleq E_i(y,y) = \lambda_+ y_+^2 - E_i^-(y_-, y_-) \triangleq \lambda_+ y_+^2 - Y_- \end{aligned}$$

and our task is to obtain an upper bound on $Y_- = E_i^-(y_-, y_-)$, which will translate through our bound on the eigenvalues of E_i away from zero into the desired bound on $|y|$.

We begin by obtaining bounds on $|E_i(g,y)|$, G_-, G, and Y. Since $|z| \ge |\mathrm{vol}(h,x,z)|$ for all z and $\mathrm{vol}(h,x,z) = \sum_j z_j F_j(h,x)$, we have

$$|E_i(g,y)| = |F_i(h,x)| \le 1.$$

Further, $\det\begin{pmatrix} E_i(g,g) & E_i(y,g) \\ E_i(g,y) & E_i(y,y) \end{pmatrix} < 0$ because E_i has signature $(1,1)$ restricted to the (y,g) plane, so by Lemma 4.3

$$Y = E_i(y,y) < \frac{R^2}{\varepsilon_2}.$$

On the other hand $-|x|_1 < \sum_j x_j F_j(x,h) = \sum_j h_j F_j(x,x)$, so

$$Y = E_i(y,y) = F_i(x,x) > -\frac{1}{h_i}\left((n-1)\frac{R^2}{\varepsilon_2}H + |x|_1\right) > -\left(\frac{nT}{\varepsilon_2}R^2 + R|x|_1\right).$$

Now $G = E_i(g,g) > 0$, being the area of the face about A_i in $D(P)$. We have $|E_i| = O(n/\varepsilon_3)$ by construction, so $G, G_- \le G + G_- \le |E_i||h|^2 = O(nH^2/\varepsilon_3)$ and similarly $G = O(nH^2/\varepsilon_3)$. On the other hand we have $G = \Omega(\varepsilon_2/R^2)$ by Lemma 4.3.

Now, observe that $\lambda_+ y_+ g_+$ is the geometric mean

$$\lambda_+ y_+ g_+ = \sqrt{(\lambda_+ g_+^2)(\lambda_+ y_+^2)} = \sqrt{(G+G_-)(Y+Y_-)}$$

and by Cauchy-Schwarz $E_i^-(y_-, g_-) \le \sqrt{G_- Y_-}$, so that

$$\begin{aligned} 1 \ge E_i(y,g) &\ge \sqrt{(G+G_-)(Y+Y_-)} - \sqrt{G_- Y_-} \\ &= \sqrt{Y_-}\frac{G}{\sqrt{G+G_-}+\sqrt{G_-}} + \sqrt{G+G_-}\frac{Y}{\sqrt{Y+Y_-}+\sqrt{Y_-}}. \end{aligned}$$

If $Y \geq 0$, it follows that

$$Y_- \leq \frac{2\sqrt{G+G_-}}{G} = O\left(\frac{n^2T^2R^2}{\varepsilon_2^2\varepsilon_3}\right).$$

If $Y < 0$, then

$$1 \geq \frac{G\sqrt{Y_-}}{2\sqrt{G+G_-}} - \frac{(-Y)\sqrt{G+G_-}}{\sqrt{Y_-}}$$

so

$$Y_- \leq \frac{2\sqrt{G+G_-}}{G}\sqrt{Y_-} + \frac{2(-Y)(G+G_-)}{G},$$

and because $X^2 \leq a + bX$ implies $X \leq \sqrt{a} + b$,

$$\sqrt{Y_-} \leq \frac{2\sqrt{G+G_-}}{G} + \frac{\sqrt{2(-Y)(G+G_-)}}{\sqrt{G}}$$

so that

$$Y_- = O\left(\max\left(\frac{G+G_-}{G^2}, (-Y)\frac{G+G_-}{G}\right)\right) = O\left(\frac{n^2T^3}{\varepsilon_2^2\varepsilon_3}R^2 + \frac{nT^2}{\varepsilon_2\varepsilon_3}R|x|_1\right).$$

In either case, using $Y \leq R/\varepsilon_2^2$ and Lemma 4.4, we have

$$|y|_2^2 = y_+^2 + |y_-|_2^2 \leq |E_i^{-1}|\left((Y+Y_-)+Y_-\right) = O\left(\frac{n^3T^3}{\varepsilon_2^2\varepsilon_3\varepsilon_4}R^2 + \frac{n^2T^2}{\varepsilon_2\varepsilon_3\varepsilon_4}R|x|_1\right)$$

and the theorem follows. □

Lemma 4.2. $|x|^2 \leq (4n/\varepsilon_3^2)\max_i |\pi_i x|^2$.

Proof. Let $i = \arg\max_i |x_i|$, and let v_j be a neighbor in T of v_i. Without loss of generality let $x_i > 0$. Then

$$(\pi_j x)_i = \frac{x_i - x_j\cos\phi_{ij}}{\sin\phi_{ij}} \geq x_i\frac{1-\cos\phi_{ij}}{\sin\phi_{ij}} = x_i\tan(\phi_{ij}/2) > x_i\phi_{ij}/2 \geq |x|_\infty\varepsilon_3/2$$

and it follows that

$$|\pi_i x| \geq |\pi_i x|_\infty > |x|_\infty\varepsilon_3/2 \geq |x|\varepsilon_3/2\sqrt{n}$$

which proves the lemma. □

Lemma 4.3. $F_i(h,h) > \varepsilon_2/R^2$.

Proof. The proof of Proposition 8 in [2] shows that a certain singular spherical polygon has angular area $\delta_i - \kappa_i$, where the singular spherical polygon is obtained by stereographic projection of each simplex of P_i^* onto a sphere of radius $1/r_i$ tangent to it. The total area of the polygon is $(\delta_i - \kappa_i)/r_i^2$ at this radius, so because projection of a plane figure onto a tangent sphere only decreases area we have $F_i(h,h) = \text{area}(P_i^*) > (\delta_i - \kappa_i)/r_i^2 > \varepsilon_2/R^2$. □

Lemma 4.4. *The inverse of the form E_i is bounded by $|E_i^{-1}| = O(n/\varepsilon_4)$.*

Proof. We follow the argument in Lemma 3.4 of [2] that the same form is non-degenerate. Let $\ell_j(y)$ be the length of the side between A_i and A_j in $D(P)$ when the altitudes h_{ij} are given by y. Since $E_i(y) = \frac{1}{2}\sum_j \ell_j(y)y_j$ it follows that $E_i(a,b) = \frac{1}{2}\sum_j \ell_j(a)b_j$. Therefore in order to bound the inverse of the form E_i it suffices to bound the inverse of the linear map ℓ.

Consider a y such that $|\ell(y)|_\infty \leq 1$; we will show $|y|_\infty = O(n/\varepsilon_4)$. Unfold the generalized polygon described by y into the plane, apex at the origin; the sides are of length $\ell_j(y)$, so the first and last vertex are a distance at most $|\ell(y)|_1 \leq n$ from each other. But the sum of the angles is at least ε_4 short of 2π, so this means all the vertices are within $O(n/\varepsilon_4)$ of the origin; and the altitudes y_j are no more than the distances from vertices to the origin, so they are also $O(n/\varepsilon_4)$ as claimed. □

The proof of Theorem 4.1 is complete.

5. Bounding the Hessian

In order to control the error in each step of our computation, we need to keep the Jacobian **J** along the whole step close to the value it started at, on which the step was based. To do this we bound the Hessian **H** when the triangulation is fixed, and we show that the Jacobian does not change discontinuously when changing radii force a new triangulation.

Each curvature κ_i is of the form $2\pi - \sum_{j,k:v_iv_jv_k\in T} \angle v_jOv_iv_k$, so in analyzing its derivatives we focus on the dihedral angles $\angle v_jOv_iv_k$. When the tetrahedron $Ov_iv_jv_k$ is embedded in $\mathbb{R}^3$, the angle $\angle v_jOv_iv_k$ is determined by elementary geometry as a smooth function of the distances among O, v_i, v_j, v_k. For a given triangulation T this makes κ a smooth function of r. Our first lemma shows that no error is introduced at the transitions where the triangulation $T(r)$ changes.

Lemma 5.1. *The Jacobian* **J** *is continuous at the boundary between radii corresponding to one triangulation and to another.*

Proof. Let r be a radius assigment consistent with more than one triangulation, say with a flat face $v_iv_jv_kv_l$ that can be triangulated by v_iv_k as $v_iv_jv_k, v_kv_lv_i$ or by v_jv_l as $v_jv_kv_l, v_lv_iv_j$. Since the Jacobian is continuous when either triangulation is fixed and r varies, it suffices to show that for neighboring radius assigments $r+\Delta r$, the curvatures κ obtained with either triangulation differ by a magnitude $O(|\Delta r|^2)$, with any coefficient determined by the polyhedral metric or the radius assigment r.

Embed the two tetrahedra $Ov_iv_jv_k, Ov_kv_lv_i$ or $Ov_jv_kv_l, Ov_lv_iv_j$ together in $\mathbb{R}^3$, with distances $[Ov_i]$, etc., taken from $r+\Delta r$. Of the ten pairwise distances between the five points in this diagram, eight are determined by M or the radii and do not vary between the v_iv_k and v_jv_l diagrams. Since the angles $\angle v_jOv_iv_k$, etc., are smooth functions of these ten distances, it suffices to show that the remaining two distances $[v_iv_k], [v_jv_l]$ differ between the diagrams by $O(|\Delta r|^2)$. Letting X denote the intersection of the geodesics v_iv_k, v_jv_l on the face $v_iv_jv_kv_l$, we have $[v_iv_k]$ in the v_iv_k diagram equal to $[v_iX] + [Xv_k]$, while in the v_jv_l diagram v_iXv_k form a triangle with the same lengths $[v_i X], [X v_k]$ and a shorter $[v_k v_k]$. The difference between $[v_i v_k]$ in the two diagrams is therefore the slack in the triangle inequality in this triangle $v_i X v_k$, which is bounded by $O(|\Delta r|^2)$ since the vertices have moved a distance $O(|\Delta r|)$ from where r placed them with v_i, X, v_k collinear. □

It now remains to control the change in $\mathbf{J}$ as r changes within any particular triangulation, which we do by bounding the Hessian.

Theorem 5.2. *The Hessian* $\mathbf{H}$ *has norm* $O\big(n^{5/2}S^{14}R^{23}/(\varepsilon_5^3 d_0^3 d_1^8 D^{14})\big)$.

Proof. It suffices to bound in absolute value each element $\frac{\partial^2 \kappa_i}{\partial r_j \partial r_k}$ of the Hessian. Since κ_i is 2π minus the sum of the dihedral angles about radius r_i, its derivatives decompose into sums of derivatives $\frac{\partial^2 \angle v_l O v_i v_m}{\partial r_j \partial r_k}$ where $v_i v_l v_m \in F(T)$. Since the geometry of each tetrahedron $O v_i v_l v_m$ is determined by its own side lengths, the only nonzero terms are where $j, k \in \{i, l, m\}$.

It therefore suffices to bound the second partial derivatives of dihedral angle AB in a tetrahedron $ABCD$ with respect to the lengths AB, AC, AD. By Lemma 5.6 below, these are degree-23 polynomials in the side lengths of $ABCD$, divided by $[ABCD]^3[ABC]^4[ABD]^4$. Since $2[ABC], 2[ABD] \geq (D/S)d_1$, $6[ABCD] \geq d_0(D/S)^2 \sin\varepsilon_5$, and each side is $O(R)$, the second derivative is $O\big(S^{14}R^{23}/(\varepsilon_5^3 d_0^3 d_1^8 D^{14})\big)$.

Now each element in the Hessian is the sum of at most n of these one-tetrahedron derivatives $\frac{\partial^2 \angle v_l O v_i v_m}{\partial r_j \partial r_k}$, and the norm of the Hessian itself is at most $n^{3/2}$ times the greatest absolute value of any of its elements, so the theorem is proved. □

Definition 5.3. For the remainder of this section, $ABCD$ is a tetrahedron and θ the dihedral angle $\angle CABD$ on AB.

Lemma 5.4.

$$\sin\theta = \frac{3}{2}\frac{[ABCD][AB]}{[ABC][ABD]}.$$

Proof. First, translate C and D parallel to AB to make BCD perpendicular to AB, which has no effect on either side of the equation. Now $[ABCD] = [BCD][AB]/3$ while $[ABC] = [BC][AB]/2$ and $[ABD] = [BD][AB]/2$, so our equation's right-hand side is $\frac{2[BCD]}{[BC][BD]} = \sin\angle CBD = \sin\theta$. □

Lemma 5.5. *Each of the derivatives* $\frac{\partial\theta}{\partial AB}, \frac{\partial\theta}{\partial AC}, \frac{\partial\theta}{\partial AD}$ *is a degree-10 polynomial in the side lengths of* $ABCD$*, divided by* $[ABCD][ABC]^2[ABD]^2$.

Proof. Write $[ABC]^2, [ABD]^2$ as polynomials in the side lengths using Heron's formula. Write $[ABCD]^2$ as a polynomial in the side lengths as follows. We have $36[ABCD]^2 = \det([\vec{AB}, \vec{AC}, \vec{AD}])^2 = \det(M)$ where $M = [\vec{AB}, \vec{AC}, \vec{AD}]^T[\vec{AB}, \vec{AC}, \vec{AD}]$. The entries of M are of the form $\vec{u}\cdot\vec{v} = \frac{1}{2}(|\vec{u}|^2 + |\vec{v}|^2 - |\vec{u}-\vec{v}|^2)$, which are polynomials in the side lengths. With Lemma 5.4, this gives $\sin^2\theta$ as a rational function of the side lengths.

Now $\frac{\partial \theta}{\partial x} = \frac{\partial \sin\theta}{\partial x}/\sqrt{1-\sin^2\theta}$ for any variable x, so the square of this first derivative is a rational function. Computing it in SAGE [13] or another computer algebra system finds that for each $x \in \{AB, AC, AD\}$, this squared derivative has numerator the square of a degree-10 polynomial with denominator $[ABCD]^2[ABC]^4[ABD]^4$. The lemma is proved. □

Lemma 5.6. *Each of the six second partial derivatives of θ in AB, AC, AD is a degree-23 polynomial in the side lengths of $ABCD$, divided by $[ABCD]^3[ABC]^4[ABD]^4$.*

Proof. By Lemma 5.5, each first partial derivative is a degree-10 polynomial divided by $[ABCD][ABC]^2[ABD]^2$. Since $[ABCD]^2, [ABC]^2, [ABD]^2$ are polynomials of degree 6, 4, 4 respectively, their logarithmic derivatives have themselves in the denominator and polynomials of degree 5, 3, 3 respectively in the numerator. The second partial derivatives therefore have an additional factor of $[ABCD]^2[ABC]^2[ABD]^2$ in the denominator and an additional degree of 13 in the numerator, proving the lemma. □

6. Intermediate Bounds

In this section we bound miscellaneous parameters in the computation in terms of the fundamental parameters $n, S, \varepsilon_1, \varepsilon_8$ and the computation-driving parameter ε_4.

6.1. Initial conditions.

Lemma 6.1. *Given a polyhedral metric space M, there exists a radius assignment r with curvature skew $\varepsilon_7 < \varepsilon_8/4\pi$, maximum radius $R = O(nD/\varepsilon_1\varepsilon_8)$, and minimum defect-curvature gap $\varepsilon_2 = \Omega(\varepsilon_1^2\varepsilon_8^3/n^2S^2)$.*

In the proof of Lemma 6.1 we require a lemma from singular spherical geometry.

Lemma 6.2. *Let C be a convex singular spherical n-gon with one interior vertex v of defect κ and each boundary vertex v_i a distance $\alpha \le vv_i \le \beta \le \pi/2$ from v. Then the perimeter $\mathrm{per}(C)$ is bounded by*

$$2\pi - \kappa - 2n(\pi/2-\alpha) \le \mathrm{per}(C) \le (2\pi-\kappa)\sin\beta.$$

Proof. Embed C in the singular spherical polygon B that results from removing a wedge of angle κ from a hemisphere.

To derive the lower bound, let the nearest point on the equator to each v_i be u_i, so that $u_iv_i \le \pi/2 - \alpha$. Then by the triangle inequality,

$$\mathrm{per}(C) = \sum_{ij} v_iv_j \ge \sum_{ij} u_iu_j - v_iu_i - u_jv_j \ge 2\pi - \kappa - 2n(\pi/2-\alpha).$$

For the upper bound, let D be the singular spherical surface obtained as the β-disk about v in B. Then C can be obtained by cutting D in turn along the geodesic extension of each of the sides of C. Each of these cuts, because it is a geodesic, is the shortest path with its winding number and is therefore shorter than the boundary it replaces, so the perimeter only decreases in this process. Therefore $\mathrm{per}(C) \leq \mathrm{per}(D) = (2\pi - \kappa)\sin\beta$. □

Proof of Lemma 6.1. Let r have the same value R on all vertices. We show that for sufficiently large $R = O(nD/\varepsilon_1\varepsilon_8)$ the assignment r is valid and satisfies the required bounds on ε_2 and ε_7. To do this it suffices to show that $\varepsilon_2 \leq \delta_i - \kappa_i \leq \varepsilon_7\varepsilon_1$ for the desired $\varepsilon_2, \varepsilon_7$ and each i.

For each vertex v_i, consider the singular spherical polygon C formed at v_i by the neighboring tetrahedra $v_iOv_jv_k$. Polygon C has one interior vertex at v_iO with defect κ_i, its perimeter is $\sum_{jk} \angle v_jv_iv_k = 2\pi - \delta_i$, and each vertex v_iv_k is convex. The spherical distance from the center v_iO to each vertex v_iv_k is $\angle Ov_iv_k = \pi/2 - \Theta(v_iv_k/R)$, which is at least $\rho_{\min} \triangleq \pi/2 - \Theta(D/R)$ and at most $\rho_{\max} \triangleq \pi/2 - \Theta(\ell/R)$. Now by Lemma 6.2 above, we have

$$2\pi - \kappa_i - 2n(\pi/2 - \rho_{\min}) \leq 2\pi - \delta_i \leq (2\pi - \kappa_i)\sin\rho_{\max}.$$

The left-hand inequality implies

$$\delta_i - \kappa_i \leq 2n(\pi/2 - \rho_{\min}) = O(nD/R)$$

so that $\delta_i - \kappa_i \leq (\varepsilon_8/4\pi)\varepsilon_1$ if $R = \Omega(nD/\varepsilon_1\varepsilon_8)$ for a sufficiently large constant factor. The right-hand inequality then implies

$$\delta_i - \kappa_i \geq (2\pi - \delta_i)\frac{1 - \sin\rho_{\max}}{\sin\rho_{\max}} \geq \varepsilon_8(1 - \sin\rho_{\max}) = \Omega(\varepsilon_8\ell^2/R^2) = \Omega(\varepsilon_1^2\varepsilon_8^3/n^2S^2)$$

so that the ε_2 bound holds. □

6.2. **Two angle bounds.**

Lemma 6.3. $\varepsilon_3 > \ell d_1/R^2$.

Proof. ε_3 is the smallest angle ϕ_{ij} from the apex O between any two vertices v_iv_j. Now $v_iv_j \geq \ell$, and the altitude from O to v_iv_j is at least d_1. Therefore $\frac{1}{2}\ell d_1 \leq [Ov_iv_j] \leq \frac{1}{2}\sin\phi_{ij}R^2$, so $\phi_{ij} > \sin\phi_{ij} \geq \ell d_1/R^2$. □

Lemma 6.4. $\varepsilon_5 > \varepsilon_2/6S$.

Proof. Suppose that a surface triangle has an angle of ϵ; we want to show $\epsilon > \varepsilon_2/6S$. Let the largest angle of that triangle be $\pi - \epsilon'$. By the law of sines, $\frac{\sin\epsilon'}{\sin\epsilon} \leq S$, so $\epsilon > \sin\epsilon \geq \sin\epsilon'/S > \epsilon'/3S$ since $\epsilon' \leq 2\pi/3$ implies $\sin\epsilon'/\epsilon' > 1/3$. It therefore suffices to show that $\epsilon' \geq \varepsilon_2/2$.

Let the angle of size $\pi - \epsilon'$ be at vertex i. Embed all of the tetrahedrons around Ov_i in space so that all the faces line up except for the one corresponding to an edge e adjacent to this angle of $\pi - \epsilon'$. The two copies of e are separated by an angle of κ_i. Letting f be the other side forming this large angle, the angle between one copy of e and the copy of f is $\pi - \epsilon'$. Now the sum of all the angles around v_i is $2\pi - \delta_i$, so apply the triangle inequality for angles twice to deduce

$$\begin{aligned}
\varepsilon_2 &\le 2\pi - (2\pi - \delta_i) - \kappa_i \\
&\le 2\pi - ((\pi - \epsilon') + \angle fe') - \kappa_i \\
&= \pi + \epsilon' - \angle fe' - \kappa_i \\
&\le \pi + \epsilon' - ((\pi - \epsilon') - \kappa_i) - \kappa_i \\
&= 2\epsilon'.
\end{aligned}$$

□

6.3. **Keeping away from the surface.** In this section we prove lower bounds separating O from the surface M. Recall that d_2 is the minimum distance from O to any vertex of M, d_1 is the minimum distance to any edge of T, and d is the minimum distance from O to any point of M.

Lemma 6.5. $d_2 = \Omega\big(D\varepsilon_1\varepsilon_4\varepsilon_5^2\varepsilon_8/(nS^4)\big)$.

Proof. This is an effective version of Lemma 4.8 of [2], on whose proof this one is based.

Let $i = \arg\min_i Ov_i$, so that $Ov_i = d_2$, and suppose that $d_2 = O\big(D\varepsilon_1\varepsilon_4\varepsilon_5^2\varepsilon_8/(nS^4)\big)$ with a small constant factor. We consider the singular spherical polygon C formed at the apex O by the tetrahedra about Ov_i. First we show that C is concave or nearly concave at each of its vertices, so that it satisfies the hypothesis of Lemma 6.9. Then we apply Lemma 6.9 and use the fact that the ratios of the κ_j are within $\varepsilon_7 \le \varepsilon_8/4\pi$ of those of the δ_j to get a contradiction.

Consider a vertex of C, the ray Ov_j. Let $v_iv_jv_k, v_jv_iv_l$ be the triangles in T adjacent to v_iv_j, and embed the two tetrahedra $Ov_iv_jv_k, Ov_jv_iv_l$ in $\mathbb{R}^3$. The angle of C at Ov_j is the dihedral angle $v_kOv_jv_l$.

By convexity, the dihedral angle $v_kv_iv_jv_l$ contains O, so if O is on the same side of plane $v_kv_jv_l$ as v_i is then the dihedral angle $v_kOv_jv_l$ does not contain v_i and is a reflex angle for C. Otherwise, the distance from O to this plane is at most $Ov_i = d_2$, and we will bound the magnitude of $\angle v_kOv_jv_l$.

By Lemma 5.4,

$$\sin \angle v_kOv_jv_l = \frac{3}{2}\frac{[Ov_kv_jv_l][Ov_j]}{[Ov_jv_k][Ov_jv_l]}.$$

Now $[Ov_kv_jv_l] \le d_2[v_kv_jv_l]/3 = O(d_2D^2)$ and $[Ov_j] \le [Ov_i] + [v_iv_j] \le D + d_2$. On the other hand $[Ov_jv_k] = (1/2)[Ov_j][Ov_k]\sin\angle v_jOv_k$, and $[Ov_j], [Ov_k] \ge \ell - d_2$ while $\angle v_kv_iv_j \le \angle v_iv_kO + \angle v_kOv_j + \angle Ov_jv_i \le \angle v_kOv_j + O(d_2/D)$ so that $\angle v_jOv_k \ge \varepsilon_5 - O(d_2/D)$, so $[Ov_jv_k] = \Omega(\ell^2\varepsilon_5)$, and similarly $[Ov_jv_l]$. Therefore $\sin\angle v_kOv_jv_l = O\left(d_2D^3/(\ell^4\varepsilon_5^2)\right) = O\left(\varepsilon_1\varepsilon_4\varepsilon_8/n\right)$, and the angle of C at Ov_i is

$$\angle v_kOv_jv_l = O\left(\varepsilon_1\varepsilon_4\varepsilon_8/n\right).$$

On the other hand observe that $\operatorname{per}(C) = \sum_{jk, v_i v_j v_k \in F(T)} \angle v_j O v_k \leq \sum_{jk} (\angle v_j v_i v_k + O(d_2/D)) = 2\pi - \delta_i + O(n d_2/D)$.

Now apply Lemma 6.9 to deduce that

$$\kappa_i + O(\varepsilon_1 \varepsilon_4 \varepsilon_8) \geq \left(1 - \frac{\operatorname{per}(C)}{2\pi}\right) \sum_{j \neq i} \kappa_j \geq \left(\frac{\delta_i}{2\pi} - O(n d_2/D)\right) \sum_{j \neq i} \kappa_j$$

so that

$$\begin{aligned}
\frac{\kappa_i}{\delta_i} + O(\varepsilon_4 \varepsilon_8) &\geq \left(\frac{1}{2\pi} - O(n d_2/\varepsilon_1 D)\right) \sum_{j \neq i} \kappa_j \\
&\geq (1 + o(\varepsilon_8)) \frac{1}{2\pi} \left(\min_j \frac{\kappa_j}{\delta_j}\right) \sum_{j \neq i} \delta_j \\
&= (1 + o(\varepsilon_8)) \frac{4\pi - \delta_i}{2\pi} \left(\min_j \frac{\kappa_j}{\delta_j}\right) \\
&\geq (1 + o(\varepsilon_8))(1 + \varepsilon_8/2\pi) \left(\min_j \frac{\kappa_j}{\delta_j}\right)
\end{aligned}$$

so that since $\kappa_i/\delta_i = \Omega(\varepsilon_4)$,

$$\frac{\kappa_i}{\delta_i} \left(\min_j \frac{\kappa_j}{\delta_j}\right)^{-1} \geq (1 + O(\varepsilon_8))^{-1} (1 + \varepsilon_8/2\pi)$$

which for a small enough constant factor on d_2 and hence on the $O(\varepsilon_8)$ term makes $\varepsilon_7 > \varepsilon_8/(4\pi)$, which is a contradiction. □

Lemma 6.6. $d_1 = \Omega(\varepsilon_4^2 \varepsilon_5^2 d_2^2 / D S^2) = \Omega\big(D \varepsilon_1^2 \varepsilon_4^4 \varepsilon_5^6 \varepsilon_8^2 / (n^2 S^{10})\big)$.

Proof. This is an effective version of Lemma 4.6 of [2], on whose proof this one is based.

Let O be distance d_1 from edge $v_i v_j$, which neighbors faces $v_i v_j v_k, v_j v_i v_l \in F(T)$. Consider the spherical quadrilateral D formed at O by the two tetrahedra $O v_i v_j v_k, O v_j v_i v_l$, and the singular spherical quadrilateral C formed by all the other tetrahedra. We will show the perimeter of C is nearly 2π for small d_1 and apply Lemma 6.8 to deduce a bound. This requires also upper and lower bounds on the side lengths of C and a lower bound on its exterior angles.

In triangle $O v_i v_j$, let the altitude from O have foot q; then $Oq = d_1$ while $v_i O, v_j O \geq d_2$, so $\angle v_j v_i q, \angle v_i v_j q = O(d_1/d_2)$. Also, $q v_i, q v_j \geq d_2 - d_1$, so q is at least distance $(d_2 - d_1) \sin \varepsilon_5$ from any of $v_i v_k, v_k v_j, v_j v_l, v_l v_i$, and O is at least $(d_2 - d_1) \sin \varepsilon_5 - d_1 = \Omega(d_2 \varepsilon_5)$ from each of these sides.

Now $\angle v_i O v_j = \pi - O(d_1/d_2)$ is the distance on the sphere between opposite vertices $O v_i, O v_j$ of D, so by the triangle inequality the perimeter of D is at least $2\pi - O(d_1/d_2)$. Each side of C is at least $\Omega(\varepsilon_5)$ and at most $\pi - \Omega(\varepsilon_5 d_2/D)$.

In spherical quadrilateral D, the two opposite angles $\angle v_k O v_i v_l$, $\angle v_l O v_j v_k$ are each within $O(d_1/\varepsilon_5 d_2)$ of the convex $\angle v_k v_i v_j v_l$ and therefore either reflex for D or else at least $\pi - O(d_1/\varepsilon_5 d_2)$. To bound the other two angles $\angle v_i O v_l v_j, \angle v_j O v_k v_i$, let the smaller of these be θ; then by Lemma 5.4,

$$\pi - \theta = O(\sin\theta) = O\left(\frac{(D^2 d_1)D}{(\varepsilon_5 d_2 D/S)^2}\right) = O\left(\frac{S^2 D d_1}{\varepsilon_5^2 d_2^2}\right).$$

Now there are two cases. In one case, $d_1 = \Omega(\varepsilon_4\varepsilon_5^2 d_2^2/DS^2)$. In the alternative, we find that each angle of D is at least $\pi - \varepsilon_4/2$ and each angle of C at most $\pi - \varepsilon_4/2$. In the latter case applying Lemma 6.8 to C finds that $2\pi - O(d_1/d_2) = 2\pi - \Omega(\varepsilon_4^2\varepsilon_5 d_2/D)$ so that $d_1 = \Omega(\varepsilon_4^2\varepsilon_5 d_2^2/D)$.

In either case $d_1 = \Omega(\min(\varepsilon_4\varepsilon_5^2 d_2^2/DS^2, \varepsilon_4^2\varepsilon_5 d_2^2/D)) = \Omega(\varepsilon_4^2\varepsilon_5^2 d_2^2/DS^2)$, and the bound on d_2 from Lemma 6.5 finishes the proof. □

Lemma 6.7.

$$d_0 = \Omega\left(\min\left(d_1\sqrt{\varepsilon_5}\varepsilon_4, \frac{d_1^{3/2}\varepsilon_4}{\sqrt{D}}, \frac{d_1^2\varepsilon_4}{DS^2}\right)\right) = \Omega\left(\frac{\varepsilon_1^4\varepsilon_4^9\varepsilon_5^{12}\varepsilon_8^4}{n^4 S^{22}} D\right).$$

Proof. This is an effective version of Lemma 4.5 of [2], on whose proof this one is based.

Let O be distance d_0 from triangle $v_i v_j v_k \in F(T)$. Consider the singular spherical polygon C cut out at O by all the tetrahedra other than $O v_i v_j v_k$. We show lower and upper bounds on the side lengths of C and lower bounds on its exterior angles, show the perimeter $\mathrm{per}(C)$ is near 2π for small d_0, and apply Lemma 6.8 to derive a bound.

The perimeter of C is the total angle about O on the faces of the tetrahedron $O v_i v_j v_k$, which is $2\pi - O(d_0^2/d_1^2)$. Each side of C is at least $\Omega(\varepsilon_5)$ and at most $\pi - \Omega(d_1/D)$.

Let θ be the smallest dihedral angle of $\angle v_i O v_j v_k, \angle v_j O v_k v_i, \angle v_k O v_i v_j$. Then by Lemma 5.4,

$$\pi - \theta = O(\sin\theta) = O\left(\frac{(D^2 d_0)D}{(d_1 D/S)^2}\right) = O\left(\frac{S^2 D d_0}{d_1^2}\right).$$

Now there are two cases. If $\theta \le \pi - \varepsilon_4/2$, then it follows immediately that $d_0 = \Omega(d_1^2\varepsilon_4/(S^2 D))$. Otherwise, $\theta > \pi - \varepsilon_4/2$, so the interior angles of C are more than $\varepsilon_4/2$. Applying Lemma 6.8, the perimeter $\mathrm{per}(C)$ is at most $2\pi - \Omega(\min(\varepsilon_4^2\varepsilon_5, \varepsilon_4^2 d_1/D))$, so that $d_0 = \Omega(\min(d_1\varepsilon_4\varepsilon_5^{1/2}, d_1^{3/2}\varepsilon_4 D^{-1/2}))$. The bound on d_1 from Lemma 6.6 finishes the proof. □

6.4. **Lemmas in spherical geometry.** These lemmas about singular spherical polygons and metrics are used in Subsection 6.3 above.

Lemma 6.8. *Let a convex singular spherical polygon have all exterior angles at least γ and all side lengths between c and $2\pi - c$. Then its perimeter is at most $2\pi - \Omega(\gamma^2 c)$.*

Proof. This is an effective version of Lemma 5.4 on pages 45–46 of [2], and we follow their proof. The proof in [2] shows that the perimeter is in general bounded by the perimeter in the nonsingular case. In this case consider any edge AB of the polygon, and observe that since the polygon is contained in the triangle ABC with exterior angles γ at A, B its perimeter is bounded by this triangle's perimeter. Since $c \leq AB \leq 2\pi - c$, the bound follows by straightforward spherical geometry. □

Lemma 6.9. *Let S be a singular spherical metric with vertices $\{v_i\}_i$, and let C be the singular spherical polygon consisting of the triangles about some distinguished vertex v_0. Suppose C has k convex vertices, each with an interior angle at least $\pi - \varepsilon$ for some $\varepsilon > 0$ and an exterior angle no more than π. Then*

$$\kappa_0 + 2\varepsilon k \geq \left(1 - \frac{\mathrm{per}(C)}{2\pi}\right) \sum_{i \neq 0} \kappa_i.$$

Proof. We reduce to Lemma 5.5 from [2] by induction. If $k = 0$, so that all vertices of C have interior angle at least π, then our statement is precisely theirs.

Otherwise, let v_i be a vertex of C with interior angle $\pi - \theta \in [\pi - \varepsilon, \pi)$. Draw the geodesic from v_i to v_0, and insert along this geodesic a pair of spherical triangles each with angle $\theta/2$ at v_i and angle $\kappa_0/2$ at v_0, meeting at a common vertex v_0'. The polygon C' and triangulation S' that result from adding these two triangles satisfy all the same conditions but with $k - 1$ convex vertices on C', so

$$\kappa_0' + 2\varepsilon(k-1) \geq \left(1 - \frac{\mathrm{per}(C')}{2\pi}\right) \sum_{i \neq 0} \kappa_i'.$$

Now C' and C have the same perimeter, $\kappa_0' \leq \kappa_0 + \theta \leq \kappa_0 + \varepsilon$, $\kappa_i' = \kappa_i - \theta \geq \kappa_i - \varepsilon$, and $\kappa_j' = \kappa_j$ for $j \notin \{0, i\}$, so it follows that

$$\kappa_0 + 2\varepsilon k \geq \kappa_0' + (2k-1)\varepsilon \geq \left(1 - \frac{\mathrm{per}(C')}{2\pi}\right) \sum_{i \neq 0} \kappa_i$$

as claimed. □

References

[1] A. D. Alexandrov. *Convex Polyhedra*. Springer-Verlag, Berlin, 2005.

[2] A. I. Bobenko and I. Izmestiev. Alexandrov's theorem, weighted Delaunay triangulations, and mixed volumes. *Annales de l'Institut Fourier*, 2006, math.DG/0609447.

[3] A. L. Cauchy. Sur les polygones et les polyèdres, seconde mémoire. *J. École Polytechnique*, XVIe Cahier, Tome IX:113–148, 1813. In OEuvres Complètes, IIe Sèrie, volume 1, pages 26–38, Paris, 1905.

[4] E. D. Demaine, M. L. Demaine, A. Lubiw, and J. O'Rourke. Enumerating foldings and unfoldings between polygons and polytopes. *Graphs Comb.*, 18:93–104, 2002.

[5] E. D. Demaine and J. O'Rourke. *Geometric Folding Algorithms*. Cambridge University Press, Cambridge, 2007.

[6] M. Fedorchuk and I. Pak. Rigidity and polynomial invariants of convex polytopes. *Duke Math. J.*, 129:371–404, 2005.

[7] David Glickenstein. Geometric triangulations and discrete Laplacians on manifolds, 2005, math/0508188v1.

[8] A. Lubiw and J. O'Rourke. When can a polygon fold to a polytope? Technical Report 048, Department of Computer Science, Smith College, 1996. Presented at *Am. Math. Soc.* Conf., 5 October 1996.

[9] J. S. B. Mitchell, D. M. Mount, and C. H. Papadimitriou. The discrete geodesic problem. *SIAM J. Comput.*, 16:647–668, 1987.

[10] I. Sabitov. The volume of a polyhedron as a function of its metric. *Fundam. Prikl. Mat.*, 2(4):1235–1246, 1996.

[11] I. Sabitov. The volume of a polyhedron as a function of its metric and algorithmical solution of the main problems in the metric theory of polyhedra. In *International School-Seminar Devoted to the N. V. Efimov's Memory*, pages 64–65, 1996.

[12] I. Sabitov. The volume as a metric invariant of polyhedra. *Discrete Comput. Geom.*, 20:405–425, 1998.

[13] W. Stein. *Sage: Open Source Mathematical Software (version 2.10.1)*. The Sage Group, 2008. http://www.sagemath.org.