ISSN 1364-0380 (on line) 1465-3060 (printed)

Geometry & Topology
Volume 6 (2002) 361–391
Published: 13 July 2002

Characterizing the Delaunay decompositions of compact hyperbolic surfaces

Gregory Leibon

Hinman Box 6188, Dartmouth College
Hanover NH 03755, USA

Email: `leibon@dartmouth.edu`

Abstract

Given a Delaunay decomposition of a compact hyperbolic surface, one may record the topological data of the decomposition, together with the intersection angles between the "empty disks" circumscribing the regions of the decomposition. The main result of this paper is a characterization of when a given topological decomposition and angle assignment can be realized as the data of an actual Delaunay decomposition of a hyperbolic surface.

AMS Classification numbers Primary: 52C26

Secondary: 30F10

Keywords: Delaunay triangulation, hyperbolic polyhedra, disk pattern

Proposed: Jean-Pierre Otal
Seconded: Benson Farb, Walter Neumann

Received: 28 March 2001
Revised: 8 July 2002

1 Introduction

Given a Delaunay decomposition of a compact hyperbolic surface, one may record the topological data of the decomposition, together with the intersection angles between the "empty disks" circumscribing the regions of the decomposition. The main result of this paper (Theorem 1) is a characterization of when a given topological decomposition and angle assignment can be realized as the data of an actual Delaunay decomposition of a closed hyperbolic surface. As a consequence of this characterization, we get a characterization of the convex ideal hyperbolic polyhedra associated to a compact surface with genus greater than one (Corollary 1); this result is a generalization of the convex ideal case of the Thurston–Andreev Theorem. Corollary 1 emerges naturally from Theorem 1, because the main ingredient in exploring the Delaunay decomposition is a triangulation production result (Proposition 1) whose proof relies on properties of the volume of hyperbolic polyhedra.

This paper is organized as follows. In Section 2, we formulate Theorem 1 and reduce its proof to three fundamental propositions: Propositions 1, 2, and 3. Section 3 contains the proof of Proposition 2 along with an exploration of Theorem 1's relationship to hyperbolic polyhedra. In Section 4, we exploit properties of hyperbolic volume in order to prove Proposition 1. Section 5 contains the proof of Proposition 3, a linear programming problem. Section 6 contains a discussion of consequences, generalizations, and natural questions. In particular, Section 6 contains a discussion of how we may obtain from Theorem 1 a natural polyhedral tessellation of the Teichmüller space of a compact Riemann surface with genus greater than one and at least one distinguished point.

2 The Delaunay decomposition

To begin this section we will recall how to construct the Delaunay decomposition of a compact boundaryless hyperbolic surface, G, with respect to a finite set of distinct specified points, $V = \{v_i\}_{i=1}^{|V|}$.

Delaunay construction As a first step, lift G to its universal cover, the hyperbolic plane, which we will always denote as H^2. To the inverse image of V, $\pi^{-1}(V)$, we apply Delaunay's empty sphere method, originally introduced by Delaunay in [5]. Namely, if a triple in $\pi^{-1}(V)$ lies on the boundary of a disk with interior empty of other points in $\pi^{-1}(V)$, then look at all the points in $\pi^{-1}(V)$ on this disk's boundary and take the convex hull of these points. This procedure will tile H^2 with geodesic polygons and this tiling will be invariant under G's deck group. Let us call the resulting "polygonal decomposition" (to be defined carefully in Section 6) the *Delaunay decomposition* of G with respect to V.

In this section, we will set up some preliminary topological structures that will be used keep track of the combinatorics of such decompositions. For the sake of simplicity we will assume that the topological structure underlying our decomposition is a triangulation and restrict our attention to compact boundaryless surfaces. Also, we assume that our geometry is always hyperbolic and, in particular, a geodesic triangulation will always mean a geodesic triangulation of a hyperbolic surface. In Section 6, we will carefully articulate a more general, but technically less pleasant, context into which the arguments explored in this paper will still apply. This will include dealing with the above needed "polygonal decompositions", as well as allowing our hyperbolic surfaces to have boundary, singularities, and corners.

The following notation will be used throughout. $\mathbf{T}$ will denote a triangulation and $||\mathbf{T}||$ will denote $\mathbf{T}$'s total space. $\mathbf{T}$'s cells will be referred to as follows: the set of vertices will be denoted as V, the set of the edges will be denoted by E, and the set of faces will be denoted by F. If S is a collection of cells in $\mathbf{T}$, then let $||S||$ denote the subspace of $||\mathbf{T}||$ corresponding to it. We use this notation since $|K|$ will always mean the cardinality of a set K. Also, we will abuse notation a bit and let, when appropriate, $c \in S$ mean that $||c|| \subset ||S||$.

In order to state the main theorem, we will need to carefully articulate the geometric quantities that arise in a Delaunay decomposition and, in particular, we will need to describe how to keep track of the intersection angles between neighboring "empty disks".

2.1 Characterizing the Delaunay decomposition

Given a geodesic triangulation and $t \in F$ we may embed t in H^2. The lifted vertices of t lie on the boundary of an Euclidean disk in the Poincare disk model of H^2. The intersection of this disk and the Poincare disk will be called the polygon's *circumscribing disk*. Notice, the boundary of the circumscribing disk can be intrinsically described by either a hyperbolic circle, a horocircle, or a banana (a connected component of the set of points of a fixed distance from a geodesic).

We will now carefully measure the angle between the circumscribing disks of neighboring faces. Fix a $t \in F$ that arises from a geodesic triangulation and embed t in H^2. For an $e \in t$ let γ^e be the geodesic in H^2 containing the edge corresponding to e. Let θ_t^e denote the angle between the boundary of t's circumscribing disk and γ^e as measured on the side of γ^e containing t's embedded image. For any $e \in E$ there are faces t_1 and t_2 such that $e \in t_1$ and $e \in t_2$ and we will let the *intersection angle* at e between the circumscribing disks containing $||t_1||$ and $||t_2||$ be defined as $\theta_{t_1}^e + \theta_{t_2}^e$.

In order to combinatorially keep track of these angle values, we first fix a triangulation, let $\{\psi_e\}_{e\in E}$ be a basis of $\mathbf{R}^{|E|}$, and let ψ^e denote the dual of ψ_e in $\{\psi_e\}_{e\in E}$'s dual basis. Throughout this paper, $p \in \mathbf{R}^{|E|}$ will denote a vector in such a basis and $\psi^e(p)$ will be called the *informal intersection angle complement* at e. Furthermore, let $\theta^e(p) = \pi - \psi^e(p)$ be called the *informal intersection angle* at e. Notice, associated to any geodesic triangulation we have a $p \in \mathbf{R}^{|E|}$ determined by the actual intersection angles between the circumscribing disks. In this situation p is said to be *realized* by this geodesic triangulation, which will be called p's *realization*. We will prove the following fact in Section 4.

Fact 1 *If $p \in \mathbf{R}^{|E|}$ is realized by a geodesic triangulation, then this realization is unique.*

A realization of $p \in \mathbf{R}^{|E|}$ forces a pair of necessary linear constraints on p. The most important of these constraints captures the fact that in an actual geodesic triangulation, each face has positive area, see formula 3. Call $p \in \mathbf{R}^{|E|}$ *feasible* if for any set of faces S we have, that

$$\sum_{e\in S} \theta^e(p) > \pi|S|.$$

In the context of a Delaunay decomposition, a feasibility inequality tends to an equality as vertices approach each other. The other necessary linear constraint on $p \in \mathbf{R}^{|E|}$, in order for it to be realized, is that at each vertex the hyperbolic structure is non-singular. Namely, if at every vertex v, p satisfies

$$\sum_{\{e|v\in e\}} \psi^e(p) = 2\pi,$$

then we will call p *non-singular*.

In order to prove the existence of a realization for a given $p \in \mathbf{R}^{|E|}$, we will assume that p arises not only from a geodesic triangulation but, in fact, from a Delaunay decomposition. In an actual Delaunay decomposition, it is straight forward to check that at each edge $\psi^e(p) \in (0, \pi)$, hence, we will call p *Delaunay* if $\psi^e(p) \in (0, \pi)$ for all e. The following is our main theorem.

Theorem 1 *A non-singular Delaunay $p \in \mathbf{R}^{|E|}$ is realized by a unique geodesic triangulation if and only if p is feasible.*

Context A Euclidean version of this theorem can be found in Bowditch's [1], where it is proved using techniques similar to those found in Thurston's proof of the Thurston–Andreev Theorem in [16]. As with Bowditch's result, the non-singularity condition is unnecessary and, in Section 6, we will drop it. The proof here relies on the triangulation production result, Proposition 1, stated in Section 2.2.

2.2 Triangulation Production

In a triangulation there are exactly $3|F|$ slots in which one can insert possible triangle angles, and we will identify the possible triangle angle values with the coordinates of $\mathbf{R}^{3|F|}$. These coordinates can be naturally indexed by $(e,t) \in E \times F$ by identifying the angle slot in t opposite to e with a basis vector α_e^t. Throughout this paper, $x \in \mathbf{R}^{3|F|}$ will denote a vector in this basis and α_t^e will denote the dual of α_e^t in the dual basis. In particular, $\alpha_t^e(x)$ will represent the angle in the (e,t) angle slot. Usually we will view these vectors and covectors geometrically. A vector will be expressed by placing its coefficients in a copy of the triangulation with dashed lines and a covector will contain its coefficients in a copy of the triangulation with solid lines, see figure 1. Angle slots not pictured will always be assumed to have zero as their coordinate's value. Notice, the pairing of a vector and a covector is achieved by placing the copy of the triangulation corresponding to the covector on top of the triangulation corresponding to the vector and multiplying the numbers living in the same angle slots. In figure 1, we see the covectors and vectors that we will be needing. For easy reference, we will now sum up the relevant definitions related directly to the pictured covectors and vectors.

Definition 1 Assume $x \in \mathbf{R}^{3|F|}$. For a triangle t with edges $\{e_1, e_2, e_3\}$, let $d^t(x) = \{\alpha_t^{e_1}(x), \alpha_t^{e_2}(x), \alpha_t^{e_3}(x)\}$ and call $d^t(x)$ the *angle data* associated to t. $k^t(x) = \sigma^t(x) - \pi$ will be called t's *curvature*. Notice, if $k^t(x) < 0$, then we may form an actual hyperbolic triangle with the angles in $d^t(x)$ which will be called $d^t(x)$'s *realization*. For each $e \in t$ denote the length of the edge e with respect to $d^t(x)$'s realization as $l_t^e(x)$. For an edge e sharing the triangles t_1 and t_2 we will let

$$\psi^e = \psi_{t_1}^e + \psi_{t_2}^e,$$

with ψ_t^e the covector defined in figure one. We will call ψ^e the *informal angle complement* at e. The *informal intersection angle* is defined as

$$\theta^e(x) = \pi - \psi^e(x).$$

Call x non-singular. Let an *angle system* be a point in

$$\mathbf{N} = \{x \mid k^t(x) < 0\ \forall t, \alpha_t^e(x) \in (0,\pi)\ \forall(e,t), r^v(x) = 2\pi\,\forall v\}.$$

A *conformal deformation* will be a vector in

$$\mathbf{C} = span\{w_e \mid \forall e\},$$

and we will call x and y *conformally equivalent* if $x - y \in \mathbf{C}$. Let

$$\mathbf{N}_x = (x + \mathbf{C})\bigcap \mathbf{N}$$

be called the *conformal class of* x. Let

$$\Psi\colon \mathbf{R}^{3|F|} \to \mathbf{R}^{|E|}$$

be the linear mapping given by

$$\Psi(x) = \sum_{e\in E} \psi^e(x)\psi_e.$$

We call x *Delaunay* or *non-singular* if $\Psi(x)$ is.

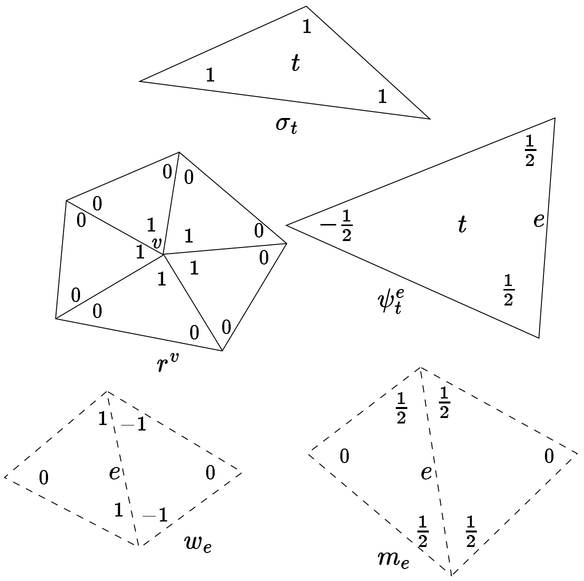

Figure 1: Here are the relevant vectors and covectors.

By the Gauss–Bonnet Theorem, we know that the curvature, $k^t(x)$, is the curvature in a geodesic triangle with angle data $d^t(x)$. Angle systems are precisely the $x \in \mathbf{R}^{3|F|}$ where the curvature is negative and all angles are realistic. In particular, the actual angle data of a geodesic triangulation is a special case of an angle system. If $u \in \mathbf{R}^{3|F|}$ is the data of a geodesic triangulation, then we will call an u *uniform*.

Our goal is to take a point in $\mathbf{N}$ and deform it into a uniform point. Such deformations are located in the affine space of conformal deformations, $\mathbf{C}$. The following is a triangulation production result whose proof (accomplished in Section 4) will utilize this conformal deformation strategy.

Proposition 1 *Every non-singular Delaunay angle system is conformally equivalent to a unique uniform angle system.*

Context The proof presented here of this triangulation production result relies on the use an objective function. It should be acknowledged that, the use of an objective function in the setting of triangulation production was first accomplished by Colin de Verdiére in [4]. Our objective function turns out to be related in a rather magical way to hyperbolic volume, and that a connection between triangulation production and hyperbolic volume should exist has its origins in Brägger's beautiful paper [2]. The particular volume exploited here was observed by my PhD thesis advisor, Peter Doyle. In the Euclidean case, this technique was carried out by Rivin in [15].

The relationship of Theorem 1 to Proposition 1 comes from the following fundamental observation, which will be proved in Section 3.

Proposition 2 *Let u be a uniform angle system and let $e \in E$, then $\theta^e(u)$ is the intersection angle of the circumscribing circles of the hyperbolic triangles sharing e.*

Proposition 2 immediately shows us that the use of the notation ψ^e and θ^e, in section 2.1, is compatible with the use in definition 1, namely, for a uniform angle system u we have that $\theta^e(u)$ really is the intersection angle θ^e.

Much of what takes place here relies on certain basic invariants of conformal deformations. For a simple example of a conformal invariant, notice that, $r^v(x)$ (from figure 1) satisfies

$$r^v(y) = r^v\left(x + \sum_{e\in E} B^e w_e\right) = r^v(x) + \sum_{e\in E} B^e r^v(w_e) = r^v(x),$$

which tells us that non-singularity is a conformal invariant. Similarly we have that

$$\psi^e(y) = \psi^e\left(x + \sum_{g\in F} B^g w_g\right) = \psi^e(x) + \sum_{g\in F} B^g \psi^e(w_g) = \psi^e(x),$$

and, hence, ψ^e and θ^e are also conformal invariants. In fact, this preservation of the formal intersection angle is why these transformations are called conformal (see [11] for a deeper reason). The next lemma expresses the fact that, the elements of $\mathbf{C}$ are the only conformal invariants in this sense.

Lemma 1 *For any $p \in \mathbf{R}^{|E|}$, $\Psi^{-1}(p) = \mathbf{C} + x$ for any $x \in \Psi^{-1}(p)$.*

Proof Since the pairing of ψ^e with the vector m_e, in figure 1, satisfies $\Psi(m_e) = \psi_e$ for each edge e, Ψ has rank E . In particular, the null space of Ψ is $3F - E = E$ dimensional. However, by the conformal invariance of ψ^e, the null space contains the E dimensional space $\mathbf{C}$; hence, the null space of Ψ is $\mathbf{C}$. □

Let $\mathbf{N}^p = \mathbf{N}_x$ for some $x \in \Psi^{-1}(p)$, and note, as a consequence of Lemma 1, Proposition 1, and Proposition 2, that a non-singular Delaunay $p \in \mathbf{R}^{|E|}$ is realized by a unique geodesic triangulation if and only if $\mathbf{N}^p$ is non-empty. Hence, Theorem 1 is now seen to be implied by the fact that $r^v(x)$ is conformally invariant and the following proposition, which is proved in Section 5.

Proposition 3 *For a Delaunay $p \in \mathbf{R}^{|E|}$, we have that $\mathbf{N}^p$ is non-empty if and only if p is feasible.*

Context Proposition 3 can be phrased as a linear programming question. In particular, since we know a solution exists, the simplex algorithm can be used to actually find the solution, see [3]. Knowing that a solution exists is usually referred to as a feasibility criteria and is the motivation for the terminology feasible. Similar geometrically motivated problems arose in [4] and [14], where they were solved using techniques from linear programming. These techniques appear to the author be much more complex to implement in this setting, hence, the optimization strategy described in Section 5.

3 Hyperbolic polyhedra: The proof of Proposition 2

Let us call H^2 the hyperbolic plane in H^3 viewed as in figure 2. The inversion, I, through the sphere of radius $\sqrt{2}$ centered at the south pole interchanges our specified H^2 with the upper half of the sphere at infinity, S_u^∞. Notice, when viewed geometrically, this map sends a point $p \in H^2$ to the point where the geodesic perpendicular to H^2 containing p hits S_u^∞ (see figure 2). In particular, being an inversion, any circle in the xy–plane is mapped to a circle on the sphere at ∞, S^∞. The use of this mapping will require the construction of an object that will be crucial in proving Proposition 1.

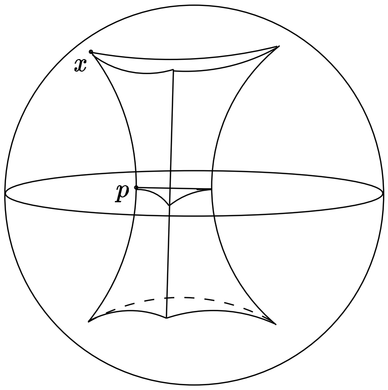

Figure 2: We have our specified hyperbolic plane $H^2 \subset H^3$ realized as the intersection of the unit sphere at the origin with the xy–plane in $\mathbf{R}^3$ via the Poincare disk model of H^3. The point p, in the figure, is mapped to the point labeled x under the inversion I. We also are viewing a triangle on that plane and its associated prism in this model.

Prism construction Place a hyperbolic triangle on a copy of $H^2 \subset H^3$. Let its associated *prism* be the convex hull of the set consisting of the triangle unioned with the geodesics perpendicular to this $H^2 \subset H^3$ going through the triangle's vertices, as visualized in figure 2. Denote the prism relative to the hyperbolic triangle constructed from the data $d = \{A, B, C\}$ as $P(d)$.

Now back to our proof. Let u be a uniform angle system, and let t_1 and t_2 be a pair of triangles sharing the edge e. Place them next to each other in our H^2 from figure 2. Notice, t_1 and t_2 have circumscribing circles in the xy–plane, which correspond to either circles, horocircles, or bananas in H^2. Since the Poincare model is conformal, the intersection angle of these circles is precisely the hyperbolic intersection angle. Being an inversion, I is conformal. In particular, these circles are sent to circles at infinity intersecting at the same angle and going through the ideal points of the neighboring $P(d_{t_1}(u))$ and $P(d_{t_2}(u))$. But these circles at infinity are also the intersection of S^∞ with the spheres representing the hyperbolic planes forming the top faces of $P(d_{t_1}(u))$ and $P(d_{t_2}(u))$. So the intersection angle of these spheres is precisely the sum of the angles inside $P(d_{t_1}(u))$ and $P(d_{t_2}(u))$ at the edge corresponding to e, which we will now see is $\theta^e(u)$. In fact, we will show that this geometric decomposition of the intersection angle is precisely the decomposition

$$\theta^e(u) = \theta^e_{t_1}(u) + \theta^e_{t_2}(u).$$

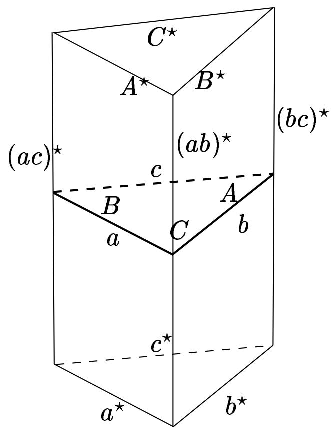

Figure 3: The notation for an ideal prism associated to hyperbolic angle data $\{A, B, C\}$, viewed for convenience in the Klein model.

Now pick an i and let $d^{t_i}(u)$ be denoted by $\{A, B, C\}$. Assume our specified edge e corresponds to the a in figure 3. From figure 4, we see the angles in figure 3 satisfy the system of linear equation telling us that interior angles of the prism sum to π at each vertex of the prism. Solving this system for the needed angle, $A^\star$, we find that indeed

$$A^\star = \frac{\pi + A - B - C}{2} = \theta^e_{t_i}(u),$$

as claimed.

We have now completed our proof of observation 2. The ideas in this proof allow us to see how a geodesic triangulation may be used to construct certain ideal hyperbolic polyhedra. Given any geodesic triangulation of $H^2 \subset H^3$, we may form the polyhedron $\bigcup_{t \in \mathbf{T}} P(d_t(u))$. We shall refer to the polyhedra that can be formed via this construction applied to the lift of a geodesic triangulation of a hyperbolic surface as the *polyhedra associated to surfaces of genus greater than 1*. As a scholium to observation 2, the dihedral angles in these polyhedra are precisely the intersection angles between the circumscribing disks of neighboring faces in the triangulation. In particular, the collection of dihedral angle complements can be naturally identified with our $p \in \mathbf{R}^{|E|}$, and results about

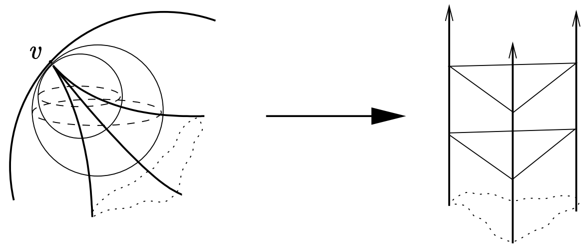

Figure 4: In this figure, we have cut off an ideal vertex of a convex hyperbolic polyhedron with a horoball. Horospheres are flat planes in H^3 and our cut produces an Euclidean polygon in this flat plane. In particular, the interior angles at an ideal vertex of a convex hyperbolic polyhedron are those of an Euclidean polygon. As indicated, this is best seen in the upper half space model of H^3, with the ideal vertex being sent to ∞. (Nice applications of this observation can be found in [16].)

geodesic triangulations can be translated into results about such polyhedra. For example, Fact 1 tells us that such polyhedra are uniquely determined by their dihedral angles, and we have the following corollary to Theorem 1 concerning the existence of such polyhedra.

Corollary 1 *A non-singular Delaunay $p \in \mathbf{R}^{|E|}$ is realizable by a unique polyhedron associated to surface of genus greater than 1 if and only if p is feasible.*

Notice, the constraint that $\theta^e(x)$ is in $(0,\pi)$, which corresponded to the Delaunay constraint, corresponds to the convex case among these polyhedra. As such, Corollary 1 can be viewed as a characterization of the possible dihedral angles in such convex polyhedra, hence, as a generalization of the convex ideal case of the Thurston–Andreev Theorem.

4 Hyperbolic volume: The proof of Proposition 1

In this section we prove Proposition 1 using an argument dependent on a rather remarkable link between hyperbolic volume and uniform structures. Let the volume of the prism $P(d^t(x))$ be denoted $V(d^t(x))$. We will be exploring the objective function

$$H(y) = \sum_{t\in\mathbf{T}} V(d^t(y)).$$

H's behavior is at the heart of our arguments. Namely, Proposition 1 follows immediately from the following claim about H's behavior, while Fact 1 follows from Lemma 1 and the following claim.

Claim 1 *View H as a function on N_x for some angle system x.*

(1) *y is a critical point of H if and only if y is uniform.*

(2) *If H has a critical point, then this critical point is unique.*

(3) *If x is Delaunay, then H has a critical point.*

Proof To prove the first part of the Claim 1 we compute H's differential. As usual, for a function on a linear space like $\mathbf{R}^{3|F|}$, we identify the tangent and cotangent spaces at every point with $\mathbf{R}^{3|F|}$ and $(\mathbf{R}^{3|F|})^*$, and express our differentials in the chosen basis. In Section 2, we prove the following lemma.

Lemma 2 *$dH(z) = \sum_{t\in F}(\sum_{e\in t} h_t^e(z)\theta_t^e)$ with the property that $h_t^e(z)$ uniquely determines the length $l_t^e(z)$ (see Definition 1).*

Recall that $T_z(\mathbf{N}_x)$ may be identified with the vector space of conformal deformations as described in definition 1. Now, simply observe that $\theta_t^f(w_e) = \pm\delta_e^f$, with the sign depending on whether t contains the negative or positive half of w_f. So, at a critical point y we have

$$0 = |dH(y)(w_e)| = \left|h_{t_1}^e(y) - h_{t_2}^e(y)\right|$$

where t_1 and t_2 are the faces sharing e. From the lemma 2 and the fact that the w_e span $T_y(\mathbf{N}_x)$, we have that $l_{t_1}^e(y) = l_{t_2}^e(y)$ for each edge is equivalent to y being critical. Now, we finish off the proof of the first part of Claim 1 by noting that, an angle system u is uniform if and only if it fits together in the sense that the lengths $l_{t_1}^e(u)$ and $l_{t_2}^e(u)$ are equal whenever t_1 and t_2 share an edge e.

In order to prove the second part of Claim 1, we need the following lemma, proved in Section 4.2.

Lemma 3 *H is a strictly concave, smooth function on $\mathbf{N}_x$ and continuous on $\mathbf{N}_x$'s closure.*

Now, a smooth concave function,like H, has at most one critical point in any open convex set, which proves part 2. To prove the third part of Claim 1, it is useful to first isolate the technical property that makes Delaunay angle systems manageable. We capture this property with the following definition. Call $x \in \partial\mathbf{N}$ *foldable* if $d^t(x) = \{0, 0, \pi\}$ for every t where either $k^t(x) = 0$ or $d^t(x) = \{A, B, 0\}$. Call a set in $\mathbf{R}^{|E|}$ *unfoldable* if the set intersects $\mathbf{N}$ but contains no foldable points. We have the following lemma.

Lemma 4 *A non-singular Delaunay angle system is unfoldable.*

Proof Let x be a non-singular Delaunay angle system. In this situation, unfoldability is quite easy to verify since $d^t(x + c)$ can never equal $\{0, 0, \pi\}$ when $x + c \in \partial\mathbf{N}$ and $\mathbf{c} \in \mathbf{C}$. To see this fact, assume to the contrary that for some t and $\mathbf{c}$ we have $d^t(x + \mathbf{c}) = \{0, 0, \pi\}$. Let e be the edge of t across from t's π and let t_1 be t's neighbor. Now notice, if $x \in \mathbf{N}$ and $d^t(x) = \{A, B, C\}$ then, since $B + C \le A + B + C = \sigma^t(x) < \pi$ and $A < \pi$, we have

$$-\frac{\pi}{2} < -\frac{A}{2} < \psi_t^e = \frac{B + C - A}{2} < \frac{B + C}{2} < \frac{\pi}{2}.$$

In other words, for any $x \in \mathbf{N}$ we have that $\psi_t^e(x) \in \left(\frac{-\pi}{2}, \frac{\pi}{2}\right)$. In particular, the conformally invariant $\psi^e(x) \in (0, \pi)$ would have to satisfy the inequality $\psi^e(x + c) = -\frac{\pi}{2} + \psi_{t_1}^e \le 0$, contradicting the fact that x is Delaunay. □

Now we will finish off the third part of Claim 1 by proving that, if $x + \mathbf{C}$ is unfoldable, then H has a critical point. To prove this, first notice, $\mathbf{N}_x$ is a pre-compact open set and H is a differentiable function on $\mathbf{N}_x$ which is continuous in $\mathbf{N}_x$'s closure, so we are assured of a critical point if H's maximum is not on $\mathbf{N}_x' s$'s boundary. Notice, since $\mathbf{N}_x$ is convex and $x + \mathbf{C}$ is unfoldable, that for every boundary point, y_0, there is a direction v such that when letting $l(s) = y_0 + sv$ we have that $l([0, \infty))$ is unfoldable. If, for some $c > 0$, we knew that

$$\lim_{s \to 0^+} \frac{d}{ds} H(l(s)) \; > \; c,$$

then there would be some $\epsilon > 0$ such that $H(l(s))$ is continuous and increasing on $[0, \epsilon)$ with $l((0, \epsilon)) \subset \mathbf{N}_x$ and, in particular, y_0 could not have been a point where$H$ achieved its maximum. This fact is immediately implied by the following lemma to be proved in Section 4.3.

Lemma 5 *For every point y_0 in $\partial\mathbf{N}$ and every direction v such that $l([0, \infty))$ is unfoldable, we have*

$$\lim_{s \to 0^+} \frac{d}{ds} H(l(s)) \; = \; \infty.$$

So, we indeed find that H achieves its unique critical in $\mathbf{N}_x$, as needed to complete the proof of the claim. □

4.1 The differential: The computation of Lemma 2

In this section we gain our needed understanding of the differential as expressed in Lemma 2. To get started, note the sum in dH is over all faces, but the fact concerns only each individual one. So we may restrict our attention to one triangle. One way to prove Lemma 2 is to explicitly compute a formula for the volume in terms of the Lobachevsky function and then find its differential. This method can be found carried out in [12]. Here we present an argument using Schlafli's formula for volume deformation. This technique has a wider range of application as well as being considerably more interesting.

To start with, we will recall Schlafli's formula for a differentiable family of compact convex polyhedra with fixed combinatorics. Let $\mathbf{E}$ denote the set of edges and $l(e)$ and $\theta(e)$ be the length and dihedral angle functions associated to an edge e. Schlafli's formula is the following formula for the deformation of the volume within this family

$$dV = -\frac{1}{2}\sum_{e\in\mathbf{E}} l(e)d(\theta(e)).$$

In the finite volume case, when there are ideal vertices, the formula changes from measuring the length of edges $l(e)$ to measuring the length of the cut off edges $l^{\star}(e)$, a fact observed by Milnor (see [15] for a proof). Let us now recall how $l^{\star}(e)$ is computed. First fix a horosphere at each ideal vertex. Then note from any horosphere to a point and between any pair of horospheres there is a unique (potentially degenerate) geodesic segment perpendicular to the horosphere(s). $l^{\star}(e)$ is the signed length of this geodesic segment; it is given a positive sign if the geodesic is outside the horosphere(s) and a negative sign if not. Schlafli's formula is independent of the horosphere choices in this construction, and I will refer to this fact as the horoball independence principle. It is worth recalling the reasoning behind this principle, since the ideas involved will come into play at several points in what follows.

The horosphere independence reasoning Recall from the proof of observation 2 that, at an ideal vertex v, we have the sum of the dihedral angles satisfying $\sum_{e\in v}\theta(e) = (n-2)\pi$, where $\{e\in v\}$ is the set of edges containing v. In particular

$$\sum_{e\in v} d\theta(e) = 0.$$

Looking at figure 4, we see by changing the horosphere at the ideal vertex v that $l^{\star}(e)$ becomes $l^{\star}(e)+c$ for each $e\in v$ with c a fixed constant. Hence, by our observation about the angle differentials

$$-2dV = \sum_{e\in\mathbf{E}} l^\star(e)d\theta(e) = \sum_{e\in\{e\in v\}^c} l^\star(e)d\theta(e) + \sum_{e\in v}(l^\star(e)+c)d\theta(e)$$

and dV is seen to be independent of the horosphere choices. □

Now let us look at our prism. Let the notation for the cut off edge lengths coincide with the edge names in figure 3. Since we may choose any horospheres, let us choose those tangent to the hyperbolic plane which our prism is symmetric across. In this case, note the lengths of $(ab)^\star$, $(bc)^\star$ and $(ac)^\star$ are zero. Recalling from the proof of observation 2 that

$$A^\star = \frac{\pi + A - B - C}{2}$$

and viewing $V(d^t(x))$ as a function on

$$\{(A,B,C)\in(0,\pi)^3 : 0 < A+B+C < \pi\}$$

we see from Schlafli's formula that

$$dV = -a^\star dA^\star - b^\star dB^\star - c^\star dC^\star.$$

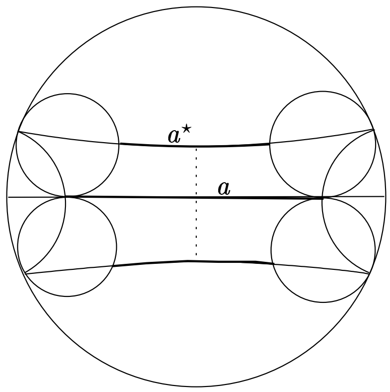

Figure 5: Here we see a face of our prism containing a. Also pictured are the horocircle slices of the horospheres tangent to the hyperbolic plane through which our prism is symmetric.

Note that Lemma 2 will follow from the following formula.

Formula 1

$$a^\star = 2\ln\left(\sinh\left(\frac{a}{2}\right)\right).$$

Proof To begin this computation, look at the face of the prism containing$a$ as in figure 5. Notice, this face is decomposed into four congruent quadrilaterals, one of which as been as in figure 6. Note that, just as with the above reasoning concerning the independence of horosphere choice, we have an independence of horocircle choice and

$$\frac{a^\star}{2} = (t^\star - h^\star) - (h^\star - s^\star).$$

In fact, $t^\star - h^\star$ and $h^\star - s^\star$ are independent of this horocircle choice as well, and it is these quantities we shall compute.

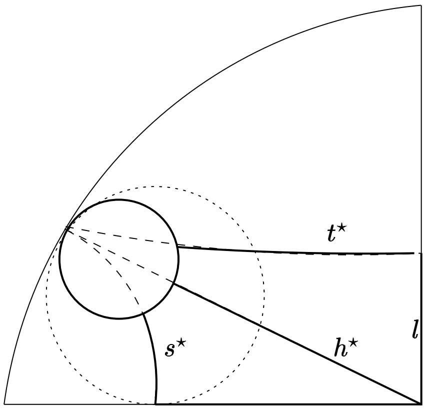

Figure 6: Pictured here is one of the four triply right-angled quadrilaterals from figure 5. Such quadrilaterals are known as Lambert quadrilaterals, and it is a well know relationship that $\tanh\left(\frac{a}{2}\right) = \mathrm{sech}(l)$, see for example [7]. (In fact it follows nearly immediately from the perhaps better known Bolyia–Lobachevsky formula.)

Look at the figure 6 and notice, using the horocircle tangent to the $\frac{a}{2}$ geodesic, that $h^\star - s^\star$ becomes precisely $h^\star$. Viewing this situation as in figure 7, we can now read off from figure 7 that

$$h^\star - s^\star = -\ln\left(\mathrm{sech}\left(\frac{a}{2}\right)\right).$$

Similarly notice, that

$$-h^\star + t^\star = \ln(\mathrm{sech}(l)),$$

which as observed in figure 6 implies

$$-h^\star + t^\star = \ln\left(\tanh\left(\frac{a}{2}\right)\right).$$

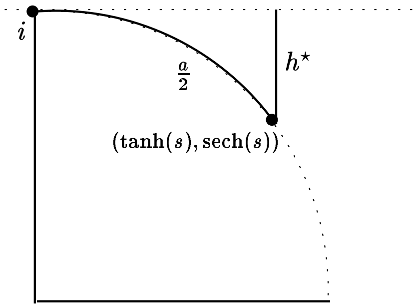

Figure 7: Here we have placed the lower triangle from figure 6 into the upper-half plane model, sending the ideal vertex to infinity and the $\frac{a}{2}$ segment on the unit circle as pictured. Recall that the unit circle in this picture can be parameterized by hyperbolic distance from i via $\tanh(s) + \operatorname{sech}(s)i$.

With these computations we now have

$$a^\star = 2\left(\ln\left(\tanh\left(\frac{a}{2}\right)\right) - \ln\left(\operatorname{sech}\left(\frac{a}{2}\right)\right)\right) = 2\ln\left(\sinh\left(\frac{a}{2}\right)\right)$$

as needed. □

4.2 Convexity: The proof of Lemma 3

To prove H is strictly concave, we start with the observation that the objective function H will certainly be a strictly concave function on $\mathbf{N}_x$ if the prism volume function $V(d^t(x))$ viewed as a function on

$$\{(A, B, C) \in (0, \pi)^3 \mid A + B + C < \pi\}$$

turned out to be strictly concave. It is worth noting that, this implies H is then strictly concave on all of $\mathbf{N}$.

There are several nice methods to explore the concavity of $V(A, B, C)$. One could simply check directly that V's Hessian is negative definite (as done in [12]), or one could exploit the visible injectivity of the gradient, or one could bootstrap from the concavity of the ideal tetrahedron's volume. It is this last method that will be presented here. The crucial observation is that any family of ideal prism can be decomposed into three ideal tetrahedra, as in figure 8. So, we have

$$V(A, B, C) = \sum_{i=1}^{3} T_i(A, B, C),$$

where T_i is the volume of the i^{th} tetrahedra in this decomposition.

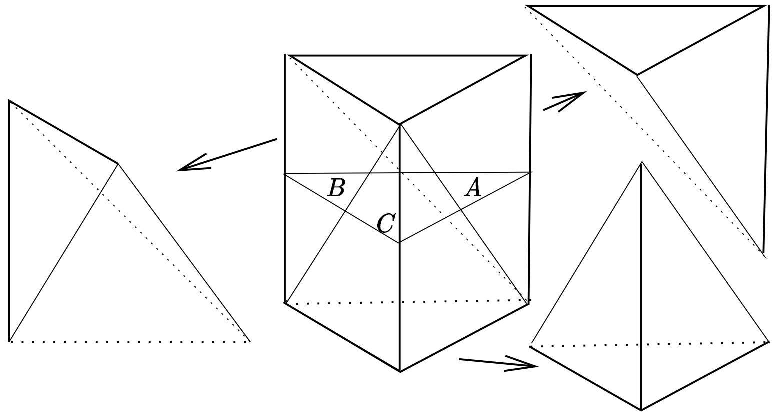

Figure 8: A decomposition of the ideal prism into three ideal tetrahedra. Notice that, the angles in this decomposition are determined by the affine conditions coming from the ideal vertices (see figure 4) along with the condition that the angles meeting along an edge slicing a prism face sum to π. In particular, all the angles depend affinely on the angles $\{A, B, C\}$.

Let us note some properties of the ideal tetrahedra and its volume. First, recall from figure 4 that the dihedral angles corresponding to the edges meeting at a vertex of an ideal tetrahedron are the angles of a Euclidean triangle. In particular, the fact that the constraint $\sum_{e\in v}\theta^e = \pi$ holds at each vertex guarantees that an ideal tetrahedron is uniquely determined by any pair of dihedral angles α and β corresponding to a pair of edges sharing a vertex. Furthermore, any pair of angles in

$$\{(\alpha, \beta) \mid \alpha + \beta < \pi\}$$

determines an ideal tetrahedron. Note the following fact (see [15]).

Fact 2 *The ideal tetrahedron's volume function, $T(\alpha, \beta)$, is strictly concave on the set*

$$\{(\alpha, \beta) \mid \alpha + \beta < \pi\}$$

and continuous on this set's closure.

From figure 8, each of the α_i and β_i of the i^{th} tetrahedron depend on the (A, B, C) affinely. This fact immediately provides us with the continuity assertion in Lemma 3. To exploit the tetrahedra's concavity we will use the following straight-forward lemma which follows immediately from the definition of concavity.

Lemma 6 *Let T be a strictly concave function on the convex set $U \subset \mathbf{R}^m$ and for each i let L_i be an affine mapping from $\mathbf{R}^n$ to $\mathbf{R}^m$ taking the convex set V into U. Then the function*

$$V(\vec{x}) = \sum_{i=1}^{k} T(L_i(\vec{x}))$$

is strictly concave on V provided $L_1 \times \ldots \times L_k$ is injective.

Letting

$$L_i(A, B, C) = (\alpha_i(A, B, C), \beta_i(A, B, C))$$

we see that V will satisfy the lemma if, for example, the mapping

$$(\alpha_1(A, B, C), \alpha_2(A, B, C)), \alpha_3(A, B, C))$$

is injective. Looking at the decomposition in figure 8, we see that we may in fact choose $\alpha_1(A, B, C) = A$, $\alpha_2(A, B, C) = B$, and $\alpha_3(A, B, C) = C$. So, indeed, we have our required injectivity and $V(A, B, C)$ is strictly concave as needed.

4.3 Boundary control: Proof of Lemma 5

To begin proving Lemma 5, note that the compactness of $\mathbf{N}$'s closure guarantees that $l(s)$ eventually hits the boundary again at some y_1 for some unique $s > 0$. So we may change the speed of our line and assume we are using the line connecting the two boundary points, namely,

$$l(s) = (1 - s)y_0 + sy_1.$$

In particular, lemma 5 is equivalent to knowing the following: for every pair of points y_0 and y_1 on $\partial\mathbf{N}$, with $l([0, 1])$ unfoldable, we have

$$\lim_{s \to 0^+} \frac{d}{ds} H(l(s)) \ = \ \infty.$$

Recalling that $H(d^t(x)) = \sum_{t \in \mathbf{T}} V(d^t(x))$, we see the lemma will follow if we demonstrate that for any triangle

$$-\infty < \lim_{s \to 0^+} \frac{d}{ds} V(d^t((s))) \leq \infty,$$

and for some triangle

$$\lim_{s \to 0^+} \frac{d}{ds} V(d^t((s))) = \infty.$$

The boundary is expressed in terms of angle data, so it would be nice to express the $-2\ln\left(\sinh\left(\frac{a}{2}\right)\right)$ coefficient in front of the $dA^\star$ term in dV (as computed in Section 4.1) in terms of the angle data. In fact, we can do even better and put this term in a form conveniently decoupling the angle and curvature.

Formula 2 $-2\ln\left(\sinh\left(\frac{a}{2}\right)\right)$ *is equal to*

$$\ln(\sin(B)) + \ln(\sin(C)) - \ln\left(\frac{\cos(A - k^t(x)) - \cos(A)}{k^t(x)}\right) - \ln(-k^t(x)).$$

Proof of formula 2 This formula relies only on the hyperbolic law of cosines which tells us

$$\cosh(a) = \frac{\cos(B)\cos(C) - \cos(A)}{\sin(A)\sin(B)}.$$

Using this relationship and the definition of $k^t(x)$ we now have

$$-2\ln\left(\sinh\left(\frac{a}{2}\right)\right) = -\ln\left(\sinh^2\left(\frac{a}{2}\right)\right) = -\ln\left(\frac{\cosh(a)-1}{2}\right)$$

$$= -\ln\left(\frac{\cos(B+C)+\cos(A)}{\sin(B)\sin(C)}\right) = -\ln\left(\frac{-\cos(A-k^t(x))+\cos(A)}{\sin(B)\sin(C)}\right),$$

as needed. □

Using this formula, we will now enumerate the possible y_0 and the behavior of $\frac{d}{ds}V(d^t((s)))$ in these various cases. Let F denote a finite constant. We will be using the fact that if $L(s)$ is an affine function of s satisfying $\lim_{s\to 0^+} L(s) = 0$ then $\lim_{s\to 0^+}\ln|\sin(L(l(s))|$ and the $\lim_{s\to 0^+}\ln|D(L(s))|$ can both be expressed as $\lim_{s\to 0^+}\ln(s) + F$. Furthermore for convenience let $d^t(y_i) = \{A_i, B_i, C_i\}$.

(1) When $d^t(y_0)$ contains no zeros and $k^t(y_0) \neq 0$, we have that

$$\lim_{s\to 0^+}\frac{d}{ds}V(d^t((s)))$$

is finite.

(2) When $d^t(y_0) = \{0, 0, \pi\}$, we have

$$\lim_{s\to 0^+}\frac{d}{ds}V(d^t((s))) =$$

$$\frac{1}{2}\lim_{s\to 0^+}\ln(s)(\sigma^t(y_1 - y_0)) - \frac{1}{2}\lim_{s\to 0^+}\ln(s)(\sigma^t(y_1 - y_0)) + F = F.$$

(3) In the case where $d^t(y_0)$ contains zeros, but $k^t(y_0) \neq 0$, for each zero (assumed to be A_0 below) we produce a term in the form

$$\lim_{s\to 0^+} \frac{d}{ds} V(d^t((s))) = \lim_{s\to 0^+} \left(-\ln(s)(A_1 - A_0)\right),$$

plus some finite quantity.

(4) When $k^t(y_0) = 0$ and no angle is zero

$$\lim_{s\to 0^+} \frac{d}{ds} V(d^t((s))) = \lim_{s\to 0^+} \frac{1}{2}\ln(s)(\sigma^t(y_1 - y_0)) + F.$$

(5) When $k^t(y_0) = 0$ and one angle, say A_0, in $d^t(y_0)$ is zero we have

$$\lim_{s\to 0^+} \frac{d}{ds} V(d^t((s))) = -2\lim_{s\to 0^+} \ln(s)(A_1 - A_0)) + \lim_{s\to 0^+} \ln(s)(\sigma^t(y_1 - y_0)) + F.$$

So, the first two cases produce finite limits. In order to understand the next three limits we make some simple observations. First, if $A_0 = 0$ and $l(s) \bigcap \mathbf{N} \neq \phi$, then $A_1 - A_0 > 0$. So, limits from the third case evaluate to $+\infty$. Secondly, note that when $k^t(y_0) = 0$ and $l(s) \bigcap \mathbf{N} \neq \phi$, that $\sigma^t(y_1 - y_0) = A_1 + B_1 + C_1 - (A_0 + B_0 + C_0) < 0$ and, hence, the limits from the fourth case are $+\infty$ too. Combining these observations we see the fifth case always produces a $+\infty$ limit as well.

For each triangle, the answer is finite or positive infinity. So, all we need to do is guarantee that for some triangle we achieve $+\infty$. To do this, note that in order for y_0 to be unfoldable, there is some triangle t such that $d^t(y_0) = \{A_0, B_0, C_0\} \neq \{0, 0, \pi\}$, however, either $k^t(y_0) = 0$ or some angle is zero. So we have at least one triangle in case 3,4, or 5, as needed.

5 Feasibility: The proof of Proposition 3

We will first demonstrate that, if $\mathbf{N}^p$ is non-empty for a Delaunay $p \in \mathbf{R}^{|E|}$, then p is feasible. In order to accomplish this, we first derive a useful formula. To state the formula, we let $O(S)$ denote the outside of S; in other words, let $O(S)$ be the set of all pairs (e, t) such that e is an edge of exactly one face in S, and t is the face in $F - S$ which contains e.

Formula 3 *For any set of triangles* S

$$\sum_{e\in S} \theta^e = \sum_{t\in S}\left(\pi - \frac{k^t}{2}\right) + \sum_{O(S)}\left(\frac{\pi}{2} - \psi_t^e\right).$$

Proof

$$\sum_E \theta^e = \sum_{e\in S}(\pi - \psi^e) = \sum_{e\in S}\left(\left(\frac{\pi}{2} - \psi^e_{t_1}\right) + \left(\frac{\pi}{2} - \psi^e_{t_2}\right)\right)$$

$$= \sum_{t\in S}\left(\pi + \frac{\pi - \sigma^t}{2}\right) + \sum_{O(S)}\left(\frac{\pi}{2} - \psi^e_t\right).$$

Substituting the definition of k^t gives the needed formula. □

Since $\mathbf{N}^p$ is non-empty, $\mathbf{N}^p$ contains a Delaunay angle system x. Apply the right and left hand sides of formula 3 to this x. Since $-k^t(x) > 0$ and $\frac{\pi}{2} - \psi^e_t(x) > 0$ (as in the proof of lemma 4), removing these terms from the right hand side strictly reduces the right hand side's size and implies that

$$\sum_{e\in S} \theta^e(x) > \pi|S|,$$

as needed.

To demonstrate the converse, that a feasible Delaunay $p \in \mathbf{R}^{|E|}$ has a non-empty $\mathbf{N}^p$, we will characterize the points in $\mathbf{N}^p$ as the minimal elements in the larger set

$$\mathbf{N}^+_p = \{x \in \mathbf{R}^{3|F|} \mid \alpha^e_t(x) \in (0,\pi), \theta^e_t(x) \in (0,\pi), \Psi(x) = p, \forall (e,t)\}$$

with respect to the objective function

$$m(x) = \left(\sum_{\{t|k_t(x)>0\}} k_t(x), |\{t : k_t(x) = 0\}|, |\{(e,t) : \alpha^t_e(x) = 0\}|\right)$$

which takes values in $\mathbf{R}\times\mathbf{Z}\times\mathbf{Z}$, given its dictionary order. By construction, $y \in \overline{\mathbf{N}^+_p}$ satisfies $m(y) = (0,0,0)$ if and only if $y \in \mathbf{N}_p$.

The key fact is that $\mathbf{N}^+_p$ is automatically non-empty, since p being Delaunay allows us to determine a point $x \in \mathbf{N}^+_p$ by setting $\theta^e_t(x) = \frac{\theta^e}{2}$. By compactness of $\overline{\mathbf{N}^+_p}$ there is a subset of $\overline{\mathbf{N}^+_p}$ upon which m achieves its minimal value. Hence, the following lemma, proved in Section 5.1, will finish our proof of Proposition 3.

Lemma 7 *The minima of m on $\overline{\mathbf{N}^+_p}$ is* $(0,0,0)$.

5.1 Minimality: The proof of Lemma 7

The proof of Lemma 7 involves the careful exploration of several "non-local" objects. This exploration, however, relies on some simple local observations, collected in the following sub-lemma.

Sub-lemma 1 *Suppose $x \in \overline{\mathbf{N}_p^+}$ for a feasible Delaunay p.*

(1) $k^t(x) < 0$, *then for every $e \in t$ we have $\theta_t^e(x) > 0$.*

(2) *Let the edges of t be $\{e, e_1, e_2\}$, $k_t(x) \geq 0$, and $\alpha_t^e(x) = 0$, then*

  (a) $\alpha_t^{e_i}(x) > 0$.

  (b) $\theta_t^{e_i}(x) > 0$.

(3) *It is impossible for pair of triangles t_1 and t_2 satisfying $k_{t_i}(x) \geq 0$ and sharing the edge e to satisfy $0 = \alpha_{t_1}^e(x) = \alpha_{t_2}^e(x)$*

(4)
$$\mathbf{N}_p^+ = \{x \in \mathbf{R}^{|E|} \mid \alpha_t^e(x) > 0, \theta_t^e(x) > 0\}$$

Proof Part 1 follows from the following fact that $\theta_t^e = \alpha_t^e - k_t^e/2$.

To begin part 2, it is useful to note that the above formula for θ_t^e demonstrates that $k^t(x) = 0$, since $\theta_t^e(x) \geq 0$. Part 2a now follows from the fact that, if there are two zeros and $k^t(x) = 0$, then the edge e where the two zeros live satisfies

$$\theta^e(x) = \theta_t^e(x) + \theta_s^e(x) = \pi + \theta_s^e(x) \geq \pi,$$

contradicting the Delaunay constraint. Similarly 2b follows by noting that if $\theta_t^{e_1}(x) = 0$, then

$$\theta_t^{e_1}(x) + \theta_t^e(x) = \pi - \alpha_t^{e_2}(x) = 0,$$

hence, $\alpha_t^{e_2}(x) = \pi$ and $\theta^{e_2}(x) \geq \pi$.

Part 3 follows by noting that, under these circumstances $\theta_{t_1}^e(x) = 0$ and $\theta_{t_2}^e(x) = 0$, hence, $\theta^e(x) = 0$, contradicting the Delaunay constraint.

Part 4 follows from the fact that, $\pi - \alpha_t^e = \theta_t^{e_1} + \theta_t^{e_2}$, where $\{e, e_1, e_2\}$ index the three edges of t, together with the fact that $\theta_{t_1}^e(x) + \theta_{t_2}^{e_2}(x) = \theta^e(x) < \pi$. □

We will now prove a weak version of Lemma 7.

Lemma 8 *The minima of m on $\overline{\mathbf{N}_p^+}$ is in the from $(0, 0, z)$.*

Proof Let x be a minimal point of m in $\overline{\mathbf{N}_p^+}$. Let Z denote the set of all triangles with $k^t(x) \geq 0$. This sub-lemma is equivalent to Z being empty. We will hypothesize that Z is not empty and produce a contradiction.

The first observation needed about Z is that $||Z||$ is not all of M. This follows from the feasibility condition on p, namely,

$$\sum_{t\in F} k^t(x) = 2(\pi|F| - \sum_{e\in S} \theta^e(x)) < 0,$$

hence, there is negative curvature somewhere.

By the minimality of x, there can be no conformal transformation capable of moving negative curvature into Z. In particular, for any $\mathbf{c} \in \mathbf{C}$ that satisfies $m(x+\mathbf{c}) < m(x)$'s, we have $x+\mathbf{c}$ is in the complement of $\overline{\mathbf{N}_p^+}$. Hence, by sub-lemma 1 part 4, we have that either $\theta_t^e(x+\mathbf{c}) < 0$ or $\alpha_t^e(x+\mathbf{c}) < 0$ for some (e,t). To exploit this fact we will attempt to conformally vacuum some curvature out of Z. Suppose we are at an internal boundary edge, e_0, of Z, ie $(e_0, t_{-1}) \in O(Z)$. Denote by t_0 the triangle in Z containing e_0. First, we will attempt to move negative curvature our of Z with a conformal transformation in the form ϵw_{e_0}. Since t_{-1} has negative curvature, by sub-lemma 1 part 1, the obstruction being able to do so must be due to a zero angle along $e_0 \in t_0$. Furthermore, since $k_{t_0}(x) \geq 0$, by sub-lemma 1 part 2a, t_0 contains only one 0.Now we will continue the attempt to remove curvature. Let e_1 be the other edge sharing the unique zero angle along e_0 in t_0; if e_1 is another boundary of Z edge we stop. If e_1 is not a boundary edge of Z use the transformation $\epsilon(w_{e_1} + w_{e_0})$ to continue the effort to remove curvature. By sub-lemma 1 part 2b, we are in the same position as above and there is a unique zero on the non t_0 side of e_1, call this face t_1. Notice, we can continue this process and form a uniquely determined set of edges and triangles, $\{e_i\}$ and $\{t_i\}$, and a sequence of conformal transformations $\epsilon\sum_{i=0}^{n} w_{e_i} \in \mathbf{C}$, see figure 9. Since such a sequence can never intersect itself with out producing a pair triangles as in the hypothesis of sub-lemma 1 part 3, the sequence must eventually poke through Z into Z^c. We will refer to the set of faces arising in this construction as a *vacuum*, see figure 9.

Now in order to utilize our vacuums, we need to remove them. Let S be the removal from Z of all these vacuums, see figure 9. The key is the following.

Sub-lemma 2 *If Z is non-empty, then S is non-empty.*

Proof First, note two distinct vacuums can never share an edge. To see this, call a vacuums *side boundary* any edge of a face in a vacuum facing a zero. Now

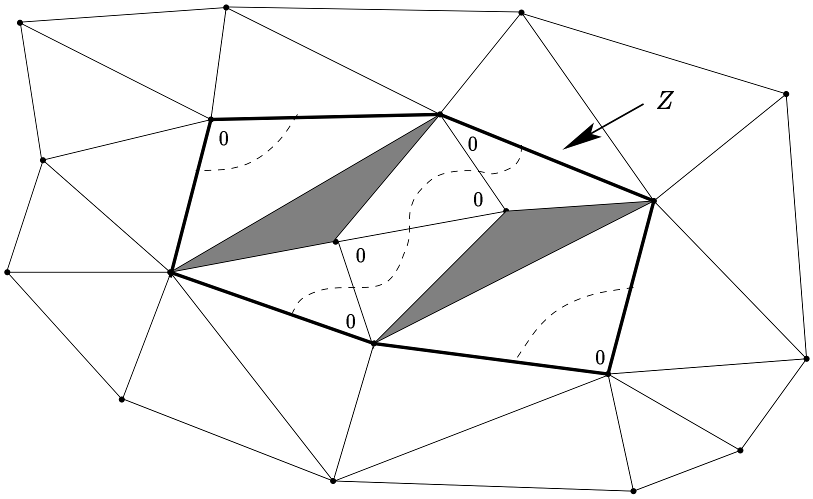

Figure 9: In this figure the region of zero curvature, $||Z||$, is the set with the bold boundary as labeled. We have also attached to Z all its vacuums and indicated them with the squiggly lines. As we can see in the figure, these vacuum transformations fail to remove curvature from Z, just as in the proof of lemma 8. Notice, the set S, Z minus its vacuums, is in this case the pair of shaded triangles.

simply note, if the intersection of two vacuums contains an edge, then it contains a first edge e_i with respect to one of the vacuums. There are two possibilities for this edge. One is that t_{i+1} has two zeros and $k^t(x) \geq 0$, contradicting sub-lemma 1 part 2a, or that e_i is a side boundary of both vacuums. In this case, we produce a pair of triangles satisfying the hypothesis of sub-lemma 1 part 3, which also can not exist.

Since vacuums cannot intersect themselves or share edges with distinct vacuums, S would be non-empty if the side boundary of a vacuum had to be in Z's interior. Look at any side boundary edge e of a fixed vacuum. Note e cannot be in Z's boundary because, in this case, the vacuum triangle it belonged to would have at least two zeros and $k^t(x) \geq 0$, contradicting sub-lemma 1 part 2a. Hence, S is nonempty. □

With our non-empty S in hand, we may now finish lemma 8 by looking at formula 3 with respect to S. Every edge in the boundary of S faces a zero on its S^c side in a triangle with $k^t(x) \geq 0$ (see figure 9 once again), so

$$\sum_{O(S)} \left(\frac{\pi}{2} - \psi_t^e(x)\right) \leq 0.$$

Also, each triangle in S has non-negative curvature, hence, from formula 3, we have violated the feasibility condition, as needed. □

We can now bootstrap from Lemma 8 to Lemma 7. Once again, assume x is a minimal point of m in $\overline{\mathbf{N}_p^+}$. From Lemma 8, we know that $m(x) = (0, 0, z)$. We will suppose this minimal z is strictly greater than zero and produce a contradiction.

Since $z > 0$ there is an angle slot which is zero. Fix such an angle slot, (e, t_0), and let $\{e, e_0, e_1\}$ index the edges of t_0. Just as in lemma 8, we now may use sub-lemma 1 part 2a to prove: in order for x not to be conformally equivalent to a point where m is decreased by the ϵw_{e_1} transformation, we must have a zero along the e_1 edge of t_0's neighbor t_1. Call this neighboring triangle t_1 and let e_2 be another edge in t_1 next to a zero angle slot in t_1. Since $k^{t_1} < 0$, we may repeat the above procedure, letting e_2 play the role of e_1 and constructing an e_3 and t_2. We may continue this process indefinitely forming a sequence of edges $\{e_i\}_{i=1}^{\infty}$. Since there are a finite number of edges, eventually the sequence $\{e_i\}_1^{\infty}$ will have some $k < l$ such that $e_k = e_l$ and $e_{k+1} = e_{l+1}$, see figure 10. I'll call such a finite set of edges an *accordion*, see in figure 10 for an example.

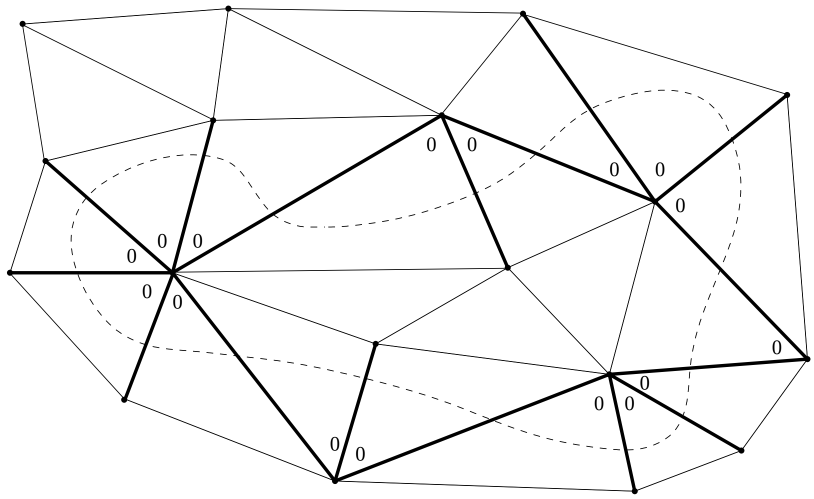

Figure 10: Here we have an accordion, the set of edges forming a loop indicated with the squiggly line. Notice, to geometrically realize the zero angles in the pictures corresponds to squeezing the accordion. The algebraic inability to squeeze various accordions is at the heart of the proof of Lemma 7.

Now we can produce a contradiction to this accordion's existence. If we denote the three edges of t_i as $\{e_i, e_{i+1}, e_{i,i+1}\}$, then we have that $\alpha_{t_i}^{e_{i,i+1}}(x) = 0$ for each i. This, together with the fact that $\alpha_{t_i}^{e_{i,i+1}} = \psi_{t_i}^{e_i} + \psi_{t_i}^{e_{i+1}}$, demonstrates that

$$0 = \sum_{i=k}^{l-1} \alpha_{t_i}^{e_{i,i+1}}(x) = \sum_{i=k}^{l-1} \psi^{e_i}(x).$$

However, by the Delaunay hypothesis on x, we have that $\sum_{i=k}^{l-1} \psi^{e_i}(x) > 0$, our needed contradiction.

6 Generalizations, consequences, and questions

6.1 Straightforward generalizations

The techniques in the previous sections apply nearly immediately to certain simple generalizations. The most important of these is to generalize the use of a triangulation to keep track of the informal angle complements. The following structure will prove much more useful. A polygonal CW–complex will be called a *polygonal decomposition* if it satisfies the following property: we may triangulate each polygon using its vertices in such way that the resulting CW–complex lifts to a triangulation of the surface's universal cover. A geodesic polygonal decomposition will be called a *Delaunay decomposition* if it can be constructed using Delaunay's empty sphere method, as described in the Delaunay Construction.

After replacing the triangulation in Section 2 with a polygonal decomposition, we have the following scholium to Theorem 1. The details of how to accomplish this can be found in [12].

Scholium 1 *Suppose our surface is compact and boundaryless. A non-singular $p \in \mathbf{R}^{|E|}$ is realized by a unique Delaunay decomposition if and only if p is feasible and Delaunay.*

An even more straight forward generalization is to to allow singularities and boundary with corners. Since the argument is localized to triangles, we may freely rid the non-singularity condition with the price that our hyperbolic surface may have specified singular points. The case of a surface with boundary follows by doubling the surface. More precisely, at a boundary edge $e \in t$ of an geodesic triangulation, we may let ψ^e be defined as ψ_t^e. If we modify the Delaunay condition to include the hypothesis that $\psi^e \in (0, \frac{\pi}{2})$ for every boundary edge e, then Theorem 1 still holds when we remove both the non-singularity hypothesis and the implicit boundaryless condition.

6.2 A consequence

We are now in a position to discuss Theorem 1's implications about the Teichmüller space of a surface with genus greater than one and at least one distinguished point. First, let us state the view of Teichmüller space that we will be exploring. Fix a differentiable surface N and a specified ordered set of points $\{v_i\}_{i=1}^{|V|}$ on N. Let (G,V,f) denote a triple consisting of a closed hyperbolic surface G, an ordered set of vertices V in G, and an orientation preserving diffeomorphism f from N to G such that $f\left(\{v_i\}_{i=1}^{|V|}\right) = V$ as an ordered set. As a set, let the Teichmüller space be the equivalence classes of triples (G,V,f), where (G,V,f) and $(\overline{G},\overline{V},\overline{f})$ will be considered equivalent if there is an isometry i from G to $\overline{G}$ such that i is isotopic to $\overline{f}\circ f^{-1}$. We will now use Scholium 1 to form a natural tessellation of Teichmüller space.

Context A combinatorially equivalent version of the tessellation of the Riemann Teichmüller space that we will construct was originally originally introduced in [8]. There have been several beautiful methods for geometrically realizing this tessellation related to our method. One of these methods starts with the wonderful observation that the Riemann Teichmüller space can be related to the Teichmüller space of singular Euclidean structures, which was first accomplished in [17]. With this view, one can then exploit the fact that the Teichmüller spaces of singular Euclidean structures has a particularly natural cellular structure, as in [1] or [18]. In fact, this was the setting where the Euclidean version of Theorem 1 was initially put to use, see Bowditch's [1]. To geometrically accomplish this identification, one first identifies the Riemann Teichmüller space with the space of complete finite area hyperbolic structures with the distinguished points viewed as cusps. Using some form of horoball decoration (as in [13]) one can relate this hyperbolic structure directly to a singular Euclidean structure using the construction in [6]. Here we will take a different geometric view in order to construct the tessellation. We will view the Riemann Teichmüller space as the space of hyperbolic structures with distinguished points. We shall find that the Delaunay decomposition and Scholium 1 provide us with the natural polyhedral cells of the Teichmüller space's tessellation. Due to the large number of excellent references, the discussion of the construction of the tessellation below will be kept relatively terse.

Tessellation construction First, we will show that the Teichmüller space can be identified, as a set, with a union of the polyhedra described in Scholium 1. We may restrict our attention to the polygonal decompositions associated to non-singular data. Let $\mathbf{P}$ be the set of such polygonal decompositions of N

viewed up to isotopies fixing $\{v_i\}_{i=1}^{|V|}$. To be explicit, the elements of $\mathbf{P}$ will be equivalence classes of pairs (P, f) where $P \in \mathbf{P}$ with an order on its vertices, V, and f is a homeomorphism of N to $||P||$ such that $f\left(\{v_i\}_{i=1}^{|V|}\right) = V$ as ordered sets. Let (P, f) and $(\bar{P}, \bar{f})$ be equivalent if there is an orientation preserving homeomorphism I of N which is isotopic to the identity relative $\{v_i\}_{i=1}^{|V|}$ and a CW equivalence h of P and $\bar{P}$, such that $I \circ f^{-1} = \bar{f}^{-1} \circ h$ when restricted to every cell. For each element $[(P, f)]$ of $\mathbf{P}$, fix an order on its edges and let $E([(P, f)])$ be the open convex polyhedron determined by the non-singular points in $(0, \pi)^{|E|}$ satisfying the feasibility condition. Notice, each point in Teichmüller space will correspond to at least one point in $E(\mathbf{P}) = \bigcup_{[(\mathbf{P},\mathbf{f})] \in \mathbf{P}} \mathbf{E}([(\mathbf{P}, \mathbf{f})])$ via the Delaunay triangulation. In fact, this correspondence is a well defined mapping, since among the decompositions in $\mathbf{P}$ there is no way to isotopically interchange the edges without interchanging the vertices as well. By the uniqueness part of Scholium 1 this mapping is injective, and the existence assertion in Scholium 1 guarantees that this mapping is surjective. So, as a set, we may identify the Teichmüller space with $E(\mathbf{P})$.

Notice that, whenever θ^e degenerates to a π, that the edge e in the underlying polygonal decomposition is removed in such a way that, after its removal, we arrive at a new polygonal decomposition with all the other data agreeing. We may use this observation to glue up $\bigcup_{[(P,f)] \in \mathbf{P}} E([(P, f)])$ in way that preserves the Euclidean structure of each $E([(P, f)])$ piece and, hence, gives $E(\mathbf{P})$ the structure of an open Euclidean cone manifold. From our set correspondence in the above construction, we have provided the Teichmüller space with a Euclidean cone manifold structure, and, in particular, the structure of a $E - V = 3\chi(N) + 2V = 6g - 6 + 2|V|$ dimensional manifold.

Now, by looking at the conformal structure of $G - V$, we have a bijective correspondence between the points in $E(\mathbf{P})$ and the Riemann Teichmüller spaceof a surface with $|V|$ punctures, and we shall call this mapping F. As a scholium to the proof of Proposition 1, as our coordinates in $E(\mathbf{P})$ continuously vary, we continuously vary the shapes of the geodesic polygons which are being used to build our hyperbolic surfaces, and in particular, are quasi-conformally deforming the conformal structure on the corresponding surfaces. So F is a bijective continuous map between $6g - 6 + 2|V|$ dimensional manifolds, hence, a homeomorphism due to the invariance of domain. Notice that, the mapping class group respects our $E([(P, f)])$ cells, and so we have formed a natural tessellation of Teichmüller space which descends to a natural tessellation of Moduli space. We sum up this discussion with the following result.

Corollary 2 *The Euclidean cone manifold $E(\mathbf{P})$ is homeomorphic to the Teichmüller space of a surface with genus greater than one and $|V|$ distinguished points. Furthermore, the Euclidean $E([(P,f)])$ cells provide a tessellation of Teichmüller space invariant under the action of the mapping class group.*

6.3 Questions

The use of hyperbolic volume in this paper to solve Proposition 1 and, hence, Theorem 1, could conceivably be used to prove the analogous questions directly in the spherical case. The prism, constructed in Section 3, becomes the following twisted prism, also observed as the right object for this game by Peter Doyle.

Twisted prism construction Fix a point p in H^3 viewed as the origin in the Poincare ball model of H^3. Imagine S^∞ is given the unit sphere's induced metric from $\mathbf{R}^3$. To construct the *twisted prism*, place a geodesic triangle on the sphere and take the convex hull of this triangle's vertices together with the point p.

To construct the associated objective function, once again, sum up the volumes of the twisted prisms associated to an angle system with positive curvature. It is easy to see that the critical points of this objective function are once again precisely the uniform angle systems, but the objective function fails to be convex. Can this objective function still be used to arrive at the analogs of Proposition 1 and Theorem 1?

A far reaching generalization of the Thurston–Andreev Theorem, due to He (see [10]), completely characterizes finite sided convex hyperbolic polyhedra. Along with the techniques in [9], this generalization could be used to characterize generalized versions of the convex polyhedra constructed here. The hyperbolic volume technique utilized here can be set up in this more general setting, though, as in the spherical case, a certain amount of control is lost. Can the hyperbolic volume techniques explored in this paper be used to prove these more general results?

References

[1] **B H Bowditch**, *Singular Euclidean structures on surfaces*, J. London Math. Soc. 44 (1991) 553–565

[2] **W Brägger**, *Kreispackungen und Triangulation*, Ens. Math. 38 (1992) 201–217

[3] **George B Dantzig**, *Linear programming and extensions*, Princeton University Press, Princeton, NJ (1963)

[4] **Y Colin de Verdiére**, *Un principe variationnel pour les empilements de cercles*, Prépublication de l'Institut Fourier, Grenoble 147 (1990) 0–16

[5] **B Delaunay**, *Sur la sphere vide*, Proc. Int. Congr. Mth. 1 (1928) 695–700

[6] **D B A Epstein**, **R C Penner**, *Euclidean decompostions of non-compact hyperbolic manifolds*, Journal of Differnetial Geometry, 27 (1988) 67–80

[7] **Marvin Jay Greenberg**, *Euclidean and non-euclidean geometries*, second edition, W H Freeman and Co. San Francisco (1972)

[8] **J L Harer**, *The virtual cohomological dimension of the mapping class group of an orientable surface*, Invent. Math. 84 (1986) 157–176

[9] **Z He**, *Rigidity of infinite disk patterns*, Ann. of Math. 149 (1999) 1–33

[10] **Zheng-Xu He**, *Disk patterns and 3–dimensional convex hyperbolic polyhedra*, in preparation

[11] **Gregory Leibon**, *Random Delaunay triangulations and metric uniformization*, IMRN, 25 (2002) 1331–1345

[12] **Gregory Leibon**, *Random Delaunay triangulations, the Thurston–Andreev Theorem and metric uniformization*, PhD thesis, UCSD (1999)

[13] **R C Penner**, *The decorated Teichmüller space of a punctured surface*, Comm. Math. Phys. 113 (1987) 299–339

[14] **I Rivin**, *Combinatorial optimization in geometry*, `arXiv:math.GT/9907032`

[15] **I Rivin**, *Euclidean structures on simplicial surfaces and hyperbolic volume*, Ann. of Math. 139 (1994) 553–580

[16] **William P Thurston**, *Three–dimensional geometry and topology*, draft edition, The Geometry Center, University of Minnesota (1991)

[17] **M Troyanov**, *Les surfaces euclidiennes a singularites coniques*, Ens. Math. 32 (1986) 74–94

[18] **W A Veech**, *Delaunay partitions*, Topology, 36 (1997) 1–28