

*Pacific
Journal of
Mathematics*

**DELAUNAY CELLS FOR ARRANGEMENTS OF FLATS
IN HYPERBOLIC SPACE**

ANDREW PRZEWORSKI

Volume 258 No. 1

July 2012

DELAUNAY CELLS FOR ARRANGEMENTS OF FLATS IN HYPERBOLIC SPACE

ANDREW PRZEWORSKI

For $n + 1$ disjoint flats of dimension k in \mathbb{H}^n , we produce a Delaunay cell that is a generalization of the Delaunay simplex associated to $n + 1$ points in \mathbb{H}^n . Combinatorially, these Delaunay cells resemble truncated n -dimensional simplices. For certain classes of arrangements of flats in \mathbb{H}^n , we prove that these Delaunay cells can be glued together to form a Delaunay complex, with the result that almost every point of \mathbb{H}^n is in a total of one Delaunay cell, counting with multiplicities and orientations.

1. Introduction

For an arrangement of points $\mathcal{P} = \{p_1, p_2, \dots\} \subset \mathbb{R}^n$ that are in general position, the Delaunay triangulation is a dual triangulation to the Voronoi tessellation. For each p_i , the corresponding Voronoi cell is the set of all points of \mathbb{R}^n that are closer to p_i than to any of the other points in \mathcal{P} . Each vertex of the Voronoi tessellation is equidistant from exactly $n + 1$ of the points of \mathcal{P} . For every vertex of the Voronoi tessellation, we can construct the convex hull of the corresponding $n + 1$ points to form a Delaunay simplex.

We generalize this to arrangements of flats in \mathbb{H}^n . Even for points, Delaunay triangulations are less well behaved in \mathbb{H}^n than in \mathbb{R}^n , so we introduce some additional assumptions.

Definition 1.1. A symmetric cocompact arrangement of k -dimensional flats in \mathbb{H}^n is a collection \mathcal{P} of k -dimensional flats for which

- (1) any two flats in \mathcal{P} have no points in common within $\overline{\mathbb{H}^n}$, and
- (2) there is a discrete group $\Gamma < \text{Isom}^+(\mathbb{H}^n)$ such that Γ acts transitively on \mathcal{P} and \mathbb{H}^n / Γ is a compact n -manifold.

Throughout this paper, \mathcal{P} will denote such an arrangement. We will use n to denote the dimension of the ambient space, k the dimension of the flats, and Γ the symmetry group. In addition, when we refer to flats being *disjoint*, we mean that they have no intersection points in $\overline{\mathbb{H}^n}$.

MSC2010: 51M09, 52C17, 57M50.

Keywords: Delaunay, hyperbolic, flats.

Given disjoint flats Π_1, \dots, Π_{n+1} , we construct a Delaunay cell. The convex hull of the flats doesn't provide a useful definition of a Delaunay cell, as the convex hull would be unbounded. Instead, for each point $\mathbf{x} \in \mathbb{H}^n$, we find the closest point $\pi_i(\mathbf{x})$ on Π_i . We define the open n -dimensional Delaunay cell to be all points $\mathbf{x} \in \mathbb{H}^n$ such that \mathbf{x} lies in the interior of the convex hull of $\pi_1(\mathbf{x}), \dots, \pi_{n+1}(\mathbf{x})$. We also define lower-dimensional open Delaunay cells to act as faces of the open n -dimensional Delaunay cell, and define a closed Delaunay cell by taking the closure of the union of the open Delaunay cells of all dimensions.

In general, the closed Delaunay cell is not a simplex, although typically its combinatorial structure resembles that of a truncated n -dimensional simplex, with the truncation faces collapsed down to the dimension of the flats. Further, the faces of the Delaunay cell are not totally geodesic surfaces. In some degenerate cases, the Delaunay cell might not even be n -dimensional. It's also possible for the Delaunay cell (or at least a part of it) to be negatively oriented. After some preliminaries in Sections 2 and 3, Section 4 is devoted to exploring the properties of these Delaunay cells.

The main result of Section 4 is that there is a surjective continuous function from a truncated n -dimensional simplex to the closed Delaunay cell and that this map can be chosen to map faces of the truncated simplex to faces of the Delaunay cell.

In Section 5, we consider the Voronoi tessellation corresponding to \mathcal{P} . For our purposes, we need for the Voronoi tessellation to have some fairly specific topological properties. We prove that for $k = n - 1$ or $n = 3, k = 1$, either the Voronoi tessellation already has the requisite properties, or it can be modified so that it does.

If we fill \mathbb{H}^n with Delaunay cells, we might hope that they would tile \mathbb{H}^n . However, that's not always the case. Some of the cells might have negative orientation and some cells might be present multiple times. In Section 6, we construct a Delaunay complex built from truncated simplices each of which maps to a Delaunay cell. We prove the following:

Corollary 6.11. *Counting cells by their multiplicities and orientations, for $n = 3, k = 1$, almost every point of \mathbb{H}^3 is in a total of one open 3-dimensional Delaunay cell.*

Theorem 6.13. *Counting cells by their multiplicities and orientations, for $k = n - 1$, almost every point of \mathbb{H}^n is in a total of one open n -dimensional Delaunay cell.*

In the specific case of $n = 2, k = 1$, Theorem 6.13 becomes an earlier result of Marshall and Martin [2003].

For $n = 3, k = 1$, the Delaunay complex also provides a geometric realization of Dehn filling. The interior of the Delaunay complex is homeomorphic

to $\mathbb{H}^3 - \bigcup_{\Pi \in \mathcal{P}} \Pi$. Taking a quotient by Γ , the Delaunay complex of truncated tetrahedra induces an ideal triangulation of $(\mathbb{H}^3 - \bigcup_{\Pi \in \mathcal{P}} \Pi) / \Gamma$, which is a cusped manifold. After Dehn filling, the cusp becomes the geodesic \mathcal{P} / Γ . In a topological sense, one can easily map the cells of the ideal triangulation of $(\mathbb{H}^3 - \bigcup_{\Pi \in \mathcal{P}} \Pi) / \Gamma$ into \mathbb{H}^3 / Γ . Mapping the Delaunay complex to Delaunay cells provides a geometric way to do this.

Although we do not explore any applications in this paper, in [Przeworski 2012], we provide upper bounds on density for packings of collars about hyperplanes in \mathbb{H}^n . One would hope that the same could be done for tube packings in compact hyperbolic 3-manifolds. The existing upper bound on the density of tube packings in \mathbb{H}^3 [Przeworski 2006] is not likely to be sharp. For ball packings, upper bounds on density are often phrased in terms of the Delaunay triangulation [Rogers 1958; Böröczky 1978].

Another possible application is to Gabai, Meyerhoff, and Milley’s MOM project [Gabai et al. 2010; 2009; Milley 2009]. MOM structures include right-angled (nonplanar) hexagons. In the case of $n = 3$, each face of one of the Delaunay cells we produce is bounded by a right-angled hexagon, and thus may provide a method of constructing a MOM structure.

2. Orthogonal projections in hyperbolic space

We use the Klein model D^n throughout for hyperbolic space, \mathbb{H}^n . In this model, hyperbolic flats are accurately represented as Euclidean flats.

Definition 2.1. We represent a k -dimensional flat Π as a vector $\mathbf{c} \in D^n$ and a set of unit vectors $\mathbf{d}_1, \dots, \mathbf{d}_k \in \mathbb{R}^n$. The vector \mathbf{c} is the point on Π that is closest to $\mathbf{0}$. The vectors $\mathbf{d}_1, \dots, \mathbf{d}_k$ are an orthonormal basis (in the Euclidean sense, not the hyperbolic sense) for the tangent space to Π . Then $\mathbf{c} \cdot \mathbf{d}_i = 0$ and $\mathbf{d}_i \cdot \mathbf{d}_j = \delta_{ij}$.

Our immediate goal is to determine the hyperbolic distance between a point $\mathbf{x} \in D^n$ and a flat Π represented as above. In order to do that, we need to be able to determine whether Euclidean vectors \mathbf{v} and \mathbf{w} in the tangent space to D^n are orthogonal in a hyperbolic sense.

Lemma 2.2 [Ratcliffe 1994]. (1) *The Riemannian metric $g_x : T_x D^n \times T_x D^n \rightarrow \mathbb{R}$ that induces hyperbolic geometry on D^n is given by*

$$g_x(\mathbf{v}, \mathbf{w}) = \frac{\mathbf{v}^T ((1 - |\mathbf{x}|^2)I_n + \mathbf{x}\mathbf{x}^T)\mathbf{w}}{(1 - |\mathbf{x}|^2)^2}.$$

(2) *The hyperbolic distance between points $\mathbf{x}, \mathbf{y} \in D^n$ is given by*

$$\cosh^{-1} \frac{1 - \mathbf{x} \cdot \mathbf{y}}{\sqrt{1 - |\mathbf{x}|^2} \sqrt{1 - |\mathbf{y}|^2}}.$$

Given a point $\mathbf{x} \in D^n$, we can now find the closest point (in a hyperbolic sense) on a flat Π .

Proposition 2.3. *Given a point $\mathbf{x} \in D^n$ and a flat Π represented by the vectors \mathbf{c} and $\mathbf{d}_1, \dots, \mathbf{d}_k$, the point on Π closest to \mathbf{x} is $\mathbf{c} + ((1 - |\mathbf{c}|^2)/(1 - \mathbf{x} \cdot \mathbf{c})) \sum_{i=1}^k (\mathbf{x} \cdot \mathbf{d}_i) \mathbf{d}_i$. We denote this point by $\pi(\mathbf{x})$, the hyperbolic projection of \mathbf{x} onto Π .*

Proof. Any point $\mathbf{y} \in \Pi$ can be written as $\mathbf{c} + \sum_{i=1}^k t_i \mathbf{d}_i$ for suitable $t_i \in \mathbb{R}$. If \mathbf{y} is to be the point on Π that is closest to \mathbf{x} , then we need $\mathbf{x} - \mathbf{y}$ to be hyperbolically orthogonal to each of $\mathbf{d}_1, \dots, \mathbf{d}_k$. Because these are vectors based at \mathbf{y} , we need $\mathbf{d}_i^T ((1 - |\mathbf{y}|^2)I_n + \mathbf{y}\mathbf{y}^T)(\mathbf{y} - \mathbf{x}) = 0$ for all i . Distributing the \mathbf{d}_i^T , this then becomes $((1 - |\mathbf{y}|^2)\mathbf{d}_i^T + t_i \mathbf{y}^T)(\mathbf{y} - \mathbf{x}) = 0$. Continuing to expand the expression, we obtain $(1 - |\mathbf{y}|^2)(t_i - \mathbf{d}_i^T \mathbf{x}) + t_i(|\mathbf{y}|^2 - \mathbf{y}^T \mathbf{x}) = 0$. Solving for t_i , we have that $t_i = ((1 - |\mathbf{y}|^2)/(1 - (\mathbf{x} \cdot \mathbf{y}))) (\mathbf{x} \cdot \mathbf{d}_i)$. That means that there is some positive number $\alpha = (1 - |\mathbf{y}|^2)/(1 - (\mathbf{x} \cdot \mathbf{y}))$ such that $t_i = \alpha(\mathbf{x} \cdot \mathbf{d}_i)$ for all i .

Using this expression for t_i , we can determine that

$$1 - |\mathbf{y}|^2 = 1 - \left(|\mathbf{c}|^2 + \sum_{i=1}^k t_i^2 \right) = 1 - \left(|\mathbf{c}|^2 + \alpha^2 \sum_{i=1}^k (\mathbf{x} \cdot \mathbf{d}_i)^2 \right).$$

Similarly,

$$1 - (\mathbf{x} \cdot \mathbf{y}) = 1 - \left((\mathbf{x} \cdot \mathbf{c}) + \alpha \sum_{i=1}^k (\mathbf{x} \cdot \mathbf{d}_i)^2 \right).$$

Using the expression $\alpha = \frac{1 - |\mathbf{y}|^2}{1 - (\mathbf{x} \cdot \mathbf{y})}$, we have that

$$\alpha \left(1 - \left((\mathbf{x} \cdot \mathbf{c}) + \alpha \sum_{i=1}^k (\mathbf{x} \cdot \mathbf{d}_i)^2 \right) \right) = 1 - \left(|\mathbf{c}|^2 + \alpha^2 \sum_{i=1}^k (\mathbf{x} \cdot \mathbf{d}_i)^2 \right).$$

Solving for α produces the result that $\alpha = \frac{1 - |\mathbf{c}|^2}{1 - \mathbf{x} \cdot \mathbf{c}}$, which proves the desired expression for $\mathbf{y} = \pi(\mathbf{x})$. \square

We sometimes regard the function π as being defined on $\overline{D^n}$.

Corollary 2.4. *If $k = n - 1$ and $\mathbf{c} \neq \mathbf{0}$, then*

$$\pi(\mathbf{x}) = \pi(\mathbf{0}) + \frac{1 - |\pi(\mathbf{0})|^2}{1 - \mathbf{x} \cdot \pi(\mathbf{0})} \left(\mathbf{x} - \frac{\mathbf{x} \cdot \pi(\mathbf{0})}{|\pi(\mathbf{0})|^2} \pi(\mathbf{0}) \right).$$

Proof. The vectors $\mathbf{c}/|\mathbf{c}|, \mathbf{d}_1, \dots, \mathbf{d}_{n-1}$ are an orthonormal basis for \mathbb{R}^n , so

$$\sum_{i=1}^{n-1} (\mathbf{x} \cdot \mathbf{d}_i) \mathbf{d}_i = \mathbf{x} - \left(\mathbf{x} \cdot \frac{\mathbf{c}}{|\mathbf{c}|} \right) \frac{\mathbf{c}}{|\mathbf{c}|}.$$

Finally, note that $\pi(\mathbf{0}) = \mathbf{c}$. \square

Lemma 2.5. *Given a flat Π in D^n , at $\mathbf{x} = \mathbf{0}$ the derivative of the projection function $\pi : D^n \rightarrow \Pi$ is equal to a scalar multiple $0 < \alpha \leq 1$ of the Euclidean orthogonal projection matrix from \mathbb{R}^n to the k -dimensional tangent space to Π . Further, $\alpha = 1$ if and only if Π contains the point $\mathbf{0}$.*

Proof. The function π is given by $\pi(\mathbf{x}) = \mathbf{c} + \frac{1-|\mathbf{c}|^2}{1-\mathbf{x} \cdot \mathbf{c}} \sum_{i=1}^k (\mathbf{x} \cdot \mathbf{d}_i) \mathbf{d}_i$. Using the product rule, the derivative is then

$$\mathbf{0} + (1 - |\mathbf{c}|^2) \left(\sum_{i=1}^k (\mathbf{x} \cdot \mathbf{d}_i) \mathbf{d}_i \frac{\mathbf{c}^T}{(1-\mathbf{x} \cdot \mathbf{c})^2} + \frac{1}{1-\mathbf{x} \cdot \mathbf{c}} \sum_{i=1}^k \mathbf{d}_i \mathbf{d}_i^T \right).$$

At $\mathbf{x} = \mathbf{0}$, this simplifies to $(1 - |\mathbf{c}|^2) \sum_{i=1}^k \mathbf{d}_i \mathbf{d}_i^T$, which proves the lemma. □

3. Vectors surrounding a point

We first recall some basic facts about barycentric coordinates.

Definition 3.1. Let points $\mathbf{v}_1, \dots, \mathbf{v}_{m+1} \in \mathbb{R}^n$ be given.

- (1) The affine hull of the points is the minimal affine subset of \mathbb{R}^n that contains all of the points.
- (2) The convex hull of the points is the minimal convex set containing all of the points.
- (3) Each point \mathbf{x} in the affine hull can be written as $\sum_{i=1}^{m+1} s_i \mathbf{v}_i$, with $\sum_{i=1}^{m+1} s_i = 1$. The s_i are called the barycentric coordinates of \mathbf{x} . If $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$ are in general position, then the barycentric coordinates of a point are unique.
- (4) If we use cyclic indexing mod $m + 1$ for the vectors \mathbf{v}_i , then the barycentric coordinates satisfy the equations

$$s_i \bigwedge_{j=i+1}^{i+m} (\mathbf{v}_j - \mathbf{v}_i) = \bigwedge_{j=i+1}^{i+m} (\mathbf{v}_j - \mathbf{x}).$$

- (5) A point \mathbf{x} is in the convex hull if and only if it has barycentric coordinates that are all nonnegative. A point is in the m -dimensional interior of the convex hull if and only if its barycentric coordinates are unique and are all positive.

Our interest is in determining which points lie in the convex hull. To do this efficiently, we need to simplify the condition a little.

Lemma 3.2. *For vectors $\mathbf{v}_1, \dots, \mathbf{v}_{m+1} \in \mathbb{R}^n$,*

$$\bigwedge_{i=1}^m (\mathbf{v}_{l+i} - \mathbf{v}_l) = \sum_{j=0}^m (-1)^{m(j-l)} \left(\bigwedge_{i=j+1}^{m+j} \mathbf{v}_i \right),$$

so the value of $(-1)^{ml} \bigwedge_{i=1}^m (\mathbf{v}_{l+i} - \mathbf{v}_l)$ is independent of l (again using cyclic indices for the vectors).

Proof. We have the expansion

$$\begin{aligned}
\bigwedge_{i=1}^m (\mathbf{v}_{l+i} - \mathbf{v}_l) &= \bigwedge_{i=1}^m \mathbf{v}_{l+i} + \sum_{j=1}^m \left(\bigwedge_{i=1}^{j-1} \mathbf{v}_{l+i} \right) \wedge (-\mathbf{v}_l) \wedge \left(\bigwedge_{i=j+1}^m \mathbf{v}_{l+i} \right) \\
&= \bigwedge_{i=1}^m \mathbf{v}_{l+i} - \sum_{j=1}^m (-1)^{j-1} \left(\bigwedge_{i=0}^{j-1} \mathbf{v}_{l+i} \right) \wedge \left(\bigwedge_{i=j+1}^m \mathbf{v}_{l+i} \right) \\
&= \bigwedge_{i=1}^m \mathbf{v}_{l+i} + \sum_{j=1}^m (-1)^j \left(\bigwedge_{i=m+1}^{m+j} \mathbf{v}_{l+i} \right) \wedge \left(\bigwedge_{i=j+1}^m \mathbf{v}_{l+i} \right) \\
&= \bigwedge_{i=1}^m \mathbf{v}_{l+i} + \sum_{j=1}^m (-1)^j (-1)^{(m-1)j} \left(\bigwedge_{i=j+1}^{m+j} \mathbf{v}_{l+i} \right) \\
&= \bigwedge_{i=1}^m \mathbf{v}_{l+i} + \sum_{j=1}^m (-1)^{mj} \left(\bigwedge_{i=j+1}^{m+j} \mathbf{v}_{l+i} \right) \\
&= \sum_{j=0}^m (-1)^{mj} \left(\bigwedge_{i=j+1}^{m+j} \mathbf{v}_{l+i} \right).
\end{aligned}$$

Reindexing i , this is the same as $\sum_{j=0}^m (-1)^{mj} \left(\bigwedge_{i=l+j+1}^{l+m+j} \mathbf{v}_i \right)$. Now, by reindexing j , this equals

$$\sum_{j=l}^{l+m} (-1)^{m(j-l)} \left(\bigwedge_{i=j+1}^{m+j} \mathbf{v}_i \right).$$

Since we are using cyclic indices mod $m+1$ and all indices are present in the j -indexed sum, we may instead sum from $j=0$ to $j=m$. Thus, we have

$$\bigwedge_{i=1}^m (\mathbf{v}_{l+i} - \mathbf{v}_l) = \sum_{j=0}^m (-1)^{m(j-l)} \left(\bigwedge_{i=j+1}^{m+j} \mathbf{v}_i \right),$$

so

$$(-1)^{ml} \bigwedge_{i=1}^m (\mathbf{v}_{l+i} - \mathbf{v}_l) = \sum_{j=0}^m (-1)^{mj} \left(\bigwedge_{i=j+1}^{m+j} \mathbf{v}_i \right).$$

Since the right side of the expression does not depend on l , it must be the case that $(-1)^{ml} \bigwedge_{i=1}^m (\mathbf{v}_{l+i} - \mathbf{v}_l)$ is also independent of l . \square

Definition 3.3. We say that a point $\mathbf{x} \in \mathbb{R}^n$ is surrounded by the points $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$ if $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$ are in general position and the point \mathbf{x} lies in the relative interior of their convex hull.

Proposition 3.4. *The point $\mathbf{x} \in \mathbb{R}^n$ is surrounded by the points $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$ (using cyclic indices mod $m + 1$) if and only if the multivectors*

$$(\ddagger) \quad (-1)^{mj} \bigwedge_{i=1}^m (\mathbf{v}_{j+i} - \mathbf{x}), \quad \text{for } j = 0, \dots, m,$$

are nonzero and are positive scalar multiples of each other.

Proof. Assume first that $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$ aren't in general position. Then \mathbf{x} is not surrounded. Further, $\bigwedge_{i=1}^m (\mathbf{v}_{1+i} - \mathbf{v}_1) = \mathbf{0}$. By Lemma 3.2,

$$\begin{aligned} \mathbf{0} &= \bigwedge_{i=1}^m (\mathbf{v}_{1+i} - \mathbf{v}_1) = \bigwedge_{i=1}^m ((\mathbf{v}_{1+i} - \mathbf{x}) - (\mathbf{v}_1 - \mathbf{x})) \\ &= \sum_{j=0}^m (-1)^{m(j-1)} \left(\bigwedge_{i=j+1}^{m+j} (\mathbf{v}_i - \mathbf{x}) \right) \\ &= (-1)^m \sum_{j=0}^m (-1)^{mj} \left(\bigwedge_{i=1}^m (\mathbf{v}_{j+i} - \mathbf{x}) \right). \end{aligned}$$

Thus, the multivectors in (\ddagger) can't all be positive scalar multiples of each other, unless they're all $\mathbf{0}$. In this case, the “if and only if” statement holds.

Now assume that $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$ are in general position. One consequence of this is that $\bigwedge_{i=j+1}^{j+m} (\mathbf{v}_i - \mathbf{v}_j)$ is nonzero.

If \mathbf{x} does not lie in the affine hull, then \mathbf{x} doesn't lie in the convex hull either. Further, if \mathbf{x} doesn't lie in the affine hull of $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$, then the affine hull of $\mathbf{x}, \mathbf{v}_1, \dots, \mathbf{v}_{m+1}$ is $(m + 1)$ -dimensional, so $\bigwedge_{i=1}^{m+1} (\mathbf{v}_{j+i} - \mathbf{x})$ is not zero. If all of the nonzero multivectors (\ddagger) are scalar multiples of each other, there is a nonzero real number α such that

$$\alpha \bigwedge_{i=1}^m (\mathbf{v}_i - \mathbf{x}) = \bigwedge_{i=1}^m (\mathbf{v}_{m+i} - \mathbf{x}).$$

Wedging both sides of the equation with $\mathbf{v}_{m+1} - \mathbf{x}$, we have $\alpha \bigwedge_{i=1}^{m+1} (\mathbf{v}_i - \mathbf{x}) = \mathbf{0}$, which is impossible. Thus, the multivectors $(-1)^{mj} \bigwedge_{i=1}^m (\mathbf{v}_{j+i} - \mathbf{x})$ are not all scalar multiples of each other, so the “if and only if” statement holds in this case.

Finally, assume that \mathbf{x} lies in the affine hull. Then \mathbf{x} is surrounded if and only if its barycentric coordinates s_j are all positive. The barycentric coordinates satisfy

$$s_j \bigwedge_{i=j+1}^{j+m} (\mathbf{v}_i - \mathbf{v}_j) = \bigwedge_{i=j+1}^{j+m} (\mathbf{v}_i - \mathbf{x}).$$

We may rewrite this as

$$(\dagger) \quad s_j(-1)^{mj} \bigwedge_{i=j+1}^{j+m} (\mathbf{v}_i - \mathbf{v}_j) = (-1)^{mj} \bigwedge_{i=j+1}^{j+m} (\mathbf{v}_i - \mathbf{x}).$$

From Lemma 3.2, we know that $(-1)^{mj} \bigwedge_{i=j+1}^{j+m} (\mathbf{v}_i - \mathbf{v}_j)$ is independent of j . Thus, \mathbf{x} is surrounded if and only if the multivectors on the right side of (\dagger) are nonzero for all j and are all positive scalar multiples of each other. These being the same multivectors as in (\ddagger) , the proposition is proved in this case too. \square

Much of this can also be done in hyperbolic geometry.

Definition 3.5. Let points $\mathbf{v}_1, \dots, \mathbf{v}_{m+1} \in \mathbb{H}^n$ be given.

- (1) The affine hull of the points is the minimal flat in \mathbb{H}^n that contains all of the points.
- (2) The convex hull of the points is the minimal convex set containing all of the points.
- (3) We say that a point $\mathbf{x} \in \mathbb{H}^n$ is surrounded by $\mathbf{p}_1, \dots, \mathbf{p}_{m+1}$ if $\mathbf{p}_1, \dots, \mathbf{p}_{m+1}$ are in general position and \mathbf{x} lies in the relative interior of their convex hull.

There are several nonequivalent ways to do define barycentric coordinates on \mathbb{H}^n . We choose the following definition.

Definition 3.6. Given points $\mathbf{p}_1, \dots, \mathbf{p}_{m+1} \in \mathbb{H}^n$ and a point \mathbf{x} in their affine hull, let $\mathbf{u}_i(\mathbf{x}) \in T_{\mathbf{x}}\mathbb{H}^n$ be the tangent vector at \mathbf{x} that is tangent to the geodesic from \mathbf{x} to \mathbf{p}_i and has length equal to the distance from \mathbf{x} to \mathbf{p}_i . Then hyperbolic barycentric coordinates of the point \mathbf{x} are $(s_1, \dots, s_{m+1}) \in \mathbb{R}^{m+1}$ such that

$$\sum_{i=1}^{m+1} s_i \mathbf{u}_i(\mathbf{x}) = \mathbf{0} \in T_{\mathbf{x}}\mathbb{H}^n \quad \text{and} \quad \sum_{i=1}^{m+1} s_i = 1.$$

Computing the hyperbolic barycentric coordinates in \mathbb{H}^n would be difficult without the use of a model. As usual, we use the Klein model, D^n .

Proposition 3.7. Given a point $\mathbf{v}_i \in D^n$ and a point $\mathbf{x} \neq \mathbf{v}_i$ in D^n , the tangent vector $\mathbf{u}_i(\mathbf{x})$ pointing from \mathbf{x} to \mathbf{v}_i (described in Definition 3.6) is

$$\frac{(1 - |\mathbf{x}|^2)(\mathbf{v}_i - \mathbf{x}) \cosh^{-1} \frac{1 - \mathbf{x} \cdot \mathbf{v}_i}{\sqrt{1 - |\mathbf{x}|^2} \sqrt{1 - |\mathbf{v}_i|^2}}}{\sqrt{(\mathbf{v}_i - \mathbf{x})^T ((1 - |\mathbf{x}|^2)I_n + \mathbf{x}\mathbf{x}^T)(\mathbf{v}_i - \mathbf{x})}}.$$

If $\mathbf{x} = \mathbf{v}_i$, then $\mathbf{u}_i(\mathbf{x}) = \mathbf{0}$. The function \mathbf{u}_i is continuously differentiable on D^n .

Proof. Identifying the set \mathbb{H}^n with the set D^n , $\mathbf{v}_i - \mathbf{x}$ is a tangent vector at \mathbf{x} that is tangent to the geodesic from \mathbf{x} to \mathbf{v}_i . Using the Riemannian metric on $T_x D^n$ given by Lemma 2.2, we see that the tangent vector $\mathbf{v}_i - \mathbf{x}$ has (hyperbolic) length

$$\sqrt{\frac{(\mathbf{v}_i - \mathbf{x})^T ((1 - |\mathbf{x}|^2)I_n + \mathbf{x}\mathbf{x}^T)(\mathbf{v}_i - \mathbf{x})}{(1 - |\mathbf{x}|^2)^2}}.$$

Then

$$\frac{(1 - |\mathbf{x}|^2)(\mathbf{v}_i - \mathbf{x}) \cosh^{-1} \frac{1 - \mathbf{x} \cdot \mathbf{v}_i}{\sqrt{1 - |\mathbf{x}|^2}\sqrt{1 - |\mathbf{v}_i|^2}}}{\sqrt{(\mathbf{v}_i - \mathbf{x})^T ((1 - |\mathbf{x}|^2)I_n + \mathbf{x}\mathbf{x}^T)(\mathbf{v}_i - \mathbf{x})}}$$

is the tangent vector $\mathbf{v}_i - \mathbf{x}$ with its (hyperbolic) length rescaled to equal the distance from \mathbf{x} to \mathbf{v}_i .

If $\mathbf{x} = \mathbf{v}_i$, then since the length of \mathbf{u}_i is equal to the distance from \mathbf{x} to \mathbf{v}_i , \mathbf{u}_i must equal $\mathbf{0}$.

The function \mathbf{u}_i was defined in an intrinsic way, so without loss of generality, we may perform a hyperbolic isometry to move \mathbf{v}_i to the origin. Hyperbolic isometries are continuously differentiable on D^n . If $\mathbf{v}_i = \mathbf{0}$ and $\mathbf{x} \neq \mathbf{0}$, then

$$\begin{aligned} \mathbf{u}_i(\mathbf{x}) &= \frac{(1 - |\mathbf{x}|^2)(-\mathbf{x}) \cosh^{-1} \frac{1}{\sqrt{1 - |\mathbf{x}|^2}}}{\sqrt{(-\mathbf{x})^T ((1 - |\mathbf{x}|^2)I_n + \mathbf{x}\mathbf{x}^T)(-\mathbf{x})}} \\ &= -\frac{(1 - |\mathbf{x}|^2)\mathbf{x} \tanh^{-1} |\mathbf{x}|}{\sqrt{|\mathbf{x}|^2 - |\mathbf{x}|^4 + |\mathbf{x}|^4}} = -(1 - |\mathbf{x}|^2)\mathbf{x} \frac{\tanh^{-1} |\mathbf{x}|}{|\mathbf{x}|}. \end{aligned}$$

Since $\mathbf{u}_i(\mathbf{0}) = \mathbf{0}$, this is a continuously differentiable function on D^n . □

We usually ignore the exceptional case where $\mathbf{x} = \mathbf{v}_i$, since the expression for $\mathbf{x} \neq \mathbf{v}_i$ has a removable singularity when $\mathbf{x} = \mathbf{v}_i$.

Proposition 3.8. *Given points $\mathbf{v}_1, \dots, \mathbf{v}_{m+1} \in D^n$ and a point \mathbf{x} in their affine hull:*

- (1) *If $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$ are in general position and the hyperbolic barycentric coordinates of \mathbf{x} exist, then the coordinates are uniquely determined.*
- (2) *The hyperbolic barycentric coordinates of \mathbf{x} exist, are unique, and are all positive if and only if \mathbf{x} is surrounded by the points $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$.*
- (3) *\mathbf{x} lies in the convex hull of the points $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$ if and only if there are hyperbolic barycentric coordinates for \mathbf{x} that are all nonnegative.*

Proof. Since the definition of the hyperbolic coordinates was given in an intrinsic manner, they must be invariant under hyperbolic isometries. Thus, without loss of

generality, we may assume that $\mathbf{x} = \mathbf{0}$. This greatly simplifies the vectors $\mathbf{u}_i(\mathbf{x})$ to

$$\mathbf{u}_i(\mathbf{0}) = \frac{\mathbf{v}_i}{|\mathbf{v}_i|} \cosh^{-1} \frac{1}{\sqrt{1 - |\mathbf{v}_i|^2}} = \mathbf{v}_i \frac{\tanh^{-1} |\mathbf{v}_i|}{|\mathbf{v}_i|}.$$

If one of the \mathbf{v}_i is $\mathbf{0}$, this expression has $\mathbf{0}$ as a limit.

The equations for computing hyperbolic barycentric coordinates are

$$\sum_{i=1}^{m+1} s_i \mathbf{v}_i \frac{\tanh^{-1} |\mathbf{v}_i|}{|\mathbf{v}_i|} = \mathbf{0} \quad \text{and} \quad \sum_{i=1}^{m+1} s_i = 1.$$

Assume that there is a solution to these equations. Let $t_i = s_i (\tanh^{-1} |\mathbf{v}_i|) / |\mathbf{v}_i|$. Then $\sum_{i=1}^{m+1} t_i \mathbf{v}_i = \mathbf{0}$. If the affine hull of $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$ contains $\mathbf{0}$ and is m -dimensional, there is a unique (up to scaling) linear combination of the vectors $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$ that adds to $\mathbf{0}$. Thus, (t_1, \dots, t_{m+1}) are determined up to scaling, so (s_1, \dots, s_{m+1}) are determined up to scaling. Since $\sum_{i=1}^{m+1} s_i = 1$, (s_1, \dots, s_{m+1}) are uniquely determined. This proves the first part of the proposition.

To prove the second part of the proposition, we use the following chain of equivalent statements:

- (1) $\mathbf{0}$ is surrounded (in a hyperbolic sense) by $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$.
- (2) $\mathbf{0}$ is surrounded (in a Euclidean sense) by $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$.
- (3) The multivectors $(-1)^{mj} \bigwedge_{i=1}^m \mathbf{v}_{j+i}$ are all nonzero and are positive scalar multiples of each other.
- (4) The multivectors $(-1)^{mj} \bigwedge_{i=1}^m \mathbf{v}_{j+i} (\tanh^{-1} |\mathbf{v}_{j+i}|) / |\mathbf{v}_{j+i}|$ are all nonzero and are positive scalar multiples of each other.
- (5) $\mathbf{0}$ is surrounded (in a Euclidean sense) by

$$\mathbf{v}_1 \frac{\tanh^{-1} |\mathbf{v}_1|}{|\mathbf{v}_1|}, \dots, \mathbf{v}_{m+1} \frac{\tanh^{-1} |\mathbf{v}_{m+1}|}{|\mathbf{v}_{m+1}|}.$$

- (6) There are uniquely determined positive numbers s_1, \dots, s_{m+1} such that

$$\sum_{i=1}^{m+1} s_i \mathbf{v}_i \frac{\tanh^{-1} |\mathbf{v}_i|}{|\mathbf{v}_i|} = \mathbf{0} \quad \text{and} \quad \sum_{i=1}^{m+1} s_i = 1.$$

- (7) $\mathbf{0}$ has uniquely determined positive hyperbolic barycentric coordinates relative to the points $\mathbf{v}_1, \dots, \mathbf{v}_{m+1} \in D^n$.

The final claim of the proposition is proved by noting that a point is on the boundary of the convex hull of $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$ if and only if either it is one of the points $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$ or it is surrounded by some proper subset of $\{\mathbf{v}_1, \dots, \mathbf{v}_{m+1}\}$. As we have already proved, a point \mathbf{x} is surrounded by some proper subset of

$\{v_1, \dots, v_{m+1}\}$ if and only if it has unique hyperbolic barycentric coordinates (relative to the proper subset) that are all positive. The remaining coordinates can all be 0. The point x is one of the v_1, \dots, v_{m+1} if and only if it has hyperbolic barycentric coordinates one of which is a 1 and the rest of which are 0. \square


4. Defining Delaunay cells

We now define the central objects of study in this paper (see figure below for $m = n = 2, k = 0, 1$).

Definition 4.1. Given pairwise disjoint k -dimensional flats Π_1, \dots, Π_{m+1} in \mathbb{H}^n (for $1 \leq m \leq n$) and the corresponding hyperbolic projection functions $\pi_i : \mathbb{H}^n \rightarrow \Pi_i$, we define

- (1) the *open m -dimensional Delaunay cell* associated to these flats to be the set of points x in \mathbb{H}^n that are surrounded by $\pi_1(x), \dots, \pi_{m+1}(x)$,
- (2) the *closed ideal m -dimensional Delaunay cell* associated to these flats to be the union of all open Delaunay cells associated with any nonsingleton subsets of the set of flats $\{\Pi_1, \dots, \Pi_{m+1}\}$, and
- (3) the *closed m -dimensional Delaunay cell* associated to these flats to be the closure in $\overline{\mathbb{H}^n}$ of the closed ideal Delaunay cell associated with the flats.

For brevity, we often omit mention of the dimension of a Delaunay cell. For purposes of intuition, we compare with the well-known Delaunay cell associated



with a set of points $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$ (that is, 0-dimensional flats). The projection function π_i would be the constant function \mathbf{v}_i . Then the open Delaunay cell would be the interior (in an m -dimensional sense) of the convex hull of $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$ (which might be empty). If nonempty, it would be an open m -dimensional simplex. The closed ideal Delaunay cell includes the interiors (in the appropriate lower-dimensional senses) of all faces of the convex hull of $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$. Assuming the open Delaunay cell is nonempty, the closed ideal Delaunay cell would be an m -dimensional simplex with the vertices removed (so topologically, though not geometrically, an ideal simplex). If the open m -dimensional cell were empty, the closed ideal m -dimensional cell would be the convex hull of $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$ with some of the vertices removed. A vertex that happened to lie in the relative interior of some other face would not be removed. Finally, in any case, the closed Delaunay cell would be the convex hull of $\mathbf{v}_1, \dots, \mathbf{v}_{m+1}$.

Proposition 4.2. *The open Delaunay cell associated with two disjoint flats Π_1 and Π_2 is the portion of their common perpendicular that lies between the two flats.*

Proof. For a point $\mathbf{x} \in D^n$ to lie in the open Delaunay cell, it would need to lie in the relative interior of the convex hull of $\pi_1(\mathbf{x})$ and $\pi_2(\mathbf{x})$. If \mathbf{x} is a point on either of the two flats, then although it lies in the convex hull of $\pi_1(\mathbf{x})$ and $\pi_2(\mathbf{x})$, it doesn't lie in the relative interior of that convex hull.

Thus, we assume that \mathbf{x} does not lie within either Π_1 or Π_2 . Let ℓ be the common perpendicular to Π_1 and Π_2 . By the definition of the projection functions, the line segment from \mathbf{x} to $\pi_i(\mathbf{x})$ must be hyperbolically perpendicular to Π_i . If \mathbf{x} were surrounded by $\pi_1(\mathbf{x})$ and $\pi_2(\mathbf{x})$, then it would lie on the open line segment joining $\pi_1(\mathbf{x})$ and $\pi_2(\mathbf{x})$. Then the points \mathbf{x} , $\pi_1(\mathbf{x})$, and $\pi_2(\mathbf{x})$ must be collinear and their containing line must be hyperbolically perpendicular to both flats. Thus, \mathbf{x} is a point on ℓ .

Since $\mathbf{x} \in \ell$, $\pi_1(\mathbf{x})$ is $\ell \cap \Pi_1$ and $\pi_2(\mathbf{x})$ is $\ell \cap \Pi_2$. The convex hull of $\pi_1(\mathbf{x})$ and $\pi_2(\mathbf{x})$ is the portion of ℓ that lies between Π_1 and Π_2 , which is a one-dimensional set that contains \mathbf{x} . Thus, $\mathbf{x} \in \ell$ is surrounded by the points $\pi_1(\mathbf{x})$ and $\pi_2(\mathbf{x})$ if and only if \mathbf{x} lies between $\Pi_1(\mathbf{x})$ and $\Pi_2(\mathbf{x})$. \square

Corollary 4.3. (1) *Closed ideal 1-dimensional Delaunay cells are open line segments.*

(2) *Closed 1-dimensional Delaunay cells are closed line segments of nonzero length.*

(3) *Closed ideal m -dimensional Delaunay cells are nonempty.*

(4) *Closed m -dimensional Delaunay cells are nonempty.*

Proof. The closed ideal 1-dimensional Delaunay cell associated with Π_1 and Π_2 is identical to the open 1-dimensional Delaunay cell associated with Π_1 and Π_2 , since

there are no (nonempty) nonsingleton subsets of $\{\Pi_1, \Pi_2\}$ except for $\{\Pi_1, \Pi_2\}$. The other claims all follow immediately. \square

Proposition 4.4. *Closed Delaunay cells are bounded subsets of \mathbb{H}^n .*

Proof. Let \mathbf{x} be a point on $\partial\mathbb{H}^n$. If \mathbf{x} is not an endpoint of any of the flats Π_1, \dots, Π_{m+1} , then the points $\pi_1(\mathbf{x}), \dots, \pi_{m+1}(\mathbf{x})$ all lie in \mathbb{H}^n , so \mathbf{x} is not in the convex hull of the points $\pi_1(\mathbf{x}), \dots, \pi_{m+1}(\mathbf{x})$. Thus, there is a neighborhood of \mathbf{x} that is disjoint from the closed Delaunay cell.

Now suppose that \mathbf{x} is an endpoint of one of the flats, say Π_1 . Then $\pi_1(\mathbf{x}) = \mathbf{x}$. Using the Klein model of \mathbb{H}^n , without loss of generality, we assume that $\mathbf{x} = (0, \dots, 0, 1)$ and that $\mathbf{0} \in \Pi_1$. The points $\pi_2(\mathbf{x}), \dots, \pi_{m+1}(\mathbf{x})$ all lie in D^n , so their n -th coordinates are strictly less than some $M < 1$. For any point $\mathbf{y} \in \overline{D^n}$, the n -th coordinate of $\pi_1(\mathbf{y})$ is equal to the n -th coordinate of \mathbf{y} . For $\mathbf{y} \in \overline{D^n}$ sufficiently near \mathbf{x} , the n -th coordinates of $\pi_2(\mathbf{y}), \dots, \pi_{m+1}(\mathbf{y})$ are all less than M , and thus less than the n -th coordinate of $\pi_1(\mathbf{y})$. Thus, \mathbf{y} does not lie in the convex hull of $\pi_1(\mathbf{y}), \dots, \pi_{m+1}(\mathbf{y})$ unless $\mathbf{y} = \pi_1(\mathbf{y})$. However, such a point would not lie in the open Delaunay cell associated to any subset of Π_1, \dots, Π_{m+1} . Thus, points sufficiently near \mathbf{x} do not lie in the closed ideal Delaunay cell, so points near \mathbf{x} don't lie in the closed ideal Delaunay cell either.

We have found a neighborhood of $\partial\mathbb{H}^n$ that is disjoint from the closed Delaunay cell. Since $\partial\mathbb{H}^n$ is homeomorphic to a sphere, it is compact. Thus, the neighborhood of $\partial\mathbb{H}^n$ can be chosen to be of uniform “thickness.” More accurately, the complement of the neighborhood can be chosen to be a closed disk of finite radius. Thus, the closed Delaunay cell is bounded. \square

We need a variant of the implicit function theorem.

Lemma 4.5. *Let $X \times Y \subset \mathbb{R}^m \times \mathbb{R}^n$ be open and connected and let $f : X \times Y \rightarrow \mathbb{R}^n$ be a continuously differentiable function denoted by $f(\mathbf{x}, \mathbf{y})$. Let $A \subset X$ be path-connected. Define $B \subset Y$ as $B = \{\mathbf{b} \in Y : \text{there is some } \mathbf{a} \in A \text{ for which } f(\mathbf{a}, \mathbf{b}) = \mathbf{0}\}$. Suppose that*

- (1) *for every $\mathbf{a} \times \mathbf{b} \in A \times B$, $\text{Det}(\partial f / \partial \mathbf{y}) \neq 0$,*
- (2) *B is bounded and its closure lies in Y , and*
- (3) *there is at least one point $\mathbf{a}_0 \in A$ for which there is a unique $\mathbf{b}_0 \in B$ such that $f(\mathbf{a}_0, \mathbf{b}_0) = \mathbf{0}$.*

Then there is a surjective differentiable function $g : A \rightarrow B$ such that $f(\mathbf{a}, \mathbf{b}) = \mathbf{0}$ if and only if $\mathbf{b} = g(\mathbf{a})$.

Proof. By the standard implicit function theorem, since $\text{Det}(\partial f / \partial \mathbf{y}) \neq 0$, if the function g exists, it is differentiable. Suppose that the function g does not exist.

Then there is some point $\mathbf{a}_1 \in A$ for which the set $\{\mathbf{b} \in B : f(\mathbf{a}_1, \mathbf{b}) = \mathbf{0}\}$ has cardinality other than 1. We refer to such a point as a point at which g is not well-defined. Let $\gamma : [0, 1] \rightarrow A$ be a path from \mathbf{a}_0 to \mathbf{a}_1 . Let

$$T = \sup\{t \in [0, 1] : g \text{ is well-defined on } \gamma([0, t])\}.$$

Without loss of generality, we may assume that either $T = 0$ or $T = 1$.

Case 1: $T = 0$. By the implicit function theorem at $\mathbf{a}_0 \times g(\mathbf{a}_0)$, for each t near 0 there is at least one point $\mathbf{b} \in B$ for which $f(\gamma(t), \mathbf{b}) = \mathbf{0}$. By the assumption that $T = 0$, there is a decreasing sequence of numbers $t_j \rightarrow 0$ such that g is not well-defined at any $\gamma(t_j)$. Thus, there must be $\mathbf{b}_j \neq \mathbf{c}_j \in B$ such that $f(\gamma(t_j), \mathbf{b}_j) = f(\gamma(t_j), \mathbf{c}_j) = \mathbf{0}$. Since B is bounded, each of these sequences has a convergent subsequence. Without loss of generality, we may assume that both sequences converge. Then

$$f(\mathbf{a}_0, \lim_{j \rightarrow \infty} \mathbf{b}_j) = \lim_{j \rightarrow \infty} f(\gamma(t_j), \mathbf{b}_j) = \mathbf{0},$$

so we must have that $\lim_{j \rightarrow \infty} \mathbf{b}_j = g(\mathbf{a}_0)$. Similarly, $\lim_{j \rightarrow \infty} \mathbf{c}_j = g(\mathbf{a}_0)$. However, by the implicit function theorem at $\mathbf{a}_0 \times g(\mathbf{a}_0)$, for each t_j there is a unique solution \mathbf{y} near $g(\mathbf{a}_0)$ to $f(\gamma(t_j), \mathbf{y}) = \mathbf{0}$. Thus, $\mathbf{b}_j = \mathbf{c}_j$, which is a contradiction.

Case 2: $T = 1$. On $\gamma([0, 1))$, g is well-defined, but at $\gamma(1) = \mathbf{a}_1$, g is not well-defined. Take an increasing sequence of numbers $t_j \rightarrow 1$. Again, since B is bounded, the sequence $g(\gamma(t_j))$ has a convergent subsequence $g(\gamma(t_{j_i}))$. Then

$$f(\mathbf{a}_1, \lim_{i \rightarrow \infty} g(\gamma(t_{j_i}))) = \lim_{i \rightarrow \infty} f(\gamma(t_{j_i}), g(\gamma(t_{j_i}))) = \mathbf{0}.$$

By the definition of the set B , $\lim_{i \rightarrow \infty} g(\gamma(t_{j_i})) \in B$. Thus, the reason that g is not well-defined at \mathbf{a}_1 must be that there are $\mathbf{b} \neq \mathbf{c} \in B$ such that $f(\mathbf{a}_1, \mathbf{b}) = f(\mathbf{a}_1, \mathbf{c}) = \mathbf{0}$. By the implicit function theorem at $\mathbf{a}_1 \times \mathbf{b}$, for large j , there is a unique \mathbf{b}_j near \mathbf{b} such that $f(\gamma(t_j), \mathbf{b}_j) = \mathbf{0}$. Similarly, for large j , there is a unique \mathbf{c}_j near \mathbf{c} such that $f(\gamma(t_j), \mathbf{c}_j) = \mathbf{0}$. Since $\mathbf{b} \neq \mathbf{c}$, we can find disjoint neighborhoods of \mathbf{b} and \mathbf{c} , so for sufficiently large j , we have $\mathbf{b}_j \neq \mathbf{c}_j$. This contradicts the assumption that g is well-defined on $\gamma([0, 1))$. □

Definition 4.6. Given pairwise disjoint flats Π_1, \dots, Π_{n+1} in D^n , define a function $f : \mathbb{R}^n \times D^n \rightarrow \mathbb{R}^n$ by using the hyperbolic barycentric coordinates equation

$$f((s_1, \dots, s_n), \mathbf{x}) = \sum_{i=1}^{n+1} \frac{(1 - |\mathbf{x}|^2)(\pi_i(\mathbf{x}) - \mathbf{x})s_i \cosh^{-1} \frac{1 - \mathbf{x} \cdot \pi_i(\mathbf{x})}{\sqrt{1 - |\mathbf{x}|^2} \sqrt{1 - |\pi_i(\mathbf{x})|^2}}}{\sqrt{(\pi_i(\mathbf{x}) - \mathbf{x})^T ((1 - |\mathbf{x}|^2)I_n + \mathbf{x}\mathbf{x}^T)(\pi_i(\mathbf{x}) - \mathbf{x})}},$$

with $s_{n+1} = 1 - \sum_{i=1}^n s_i$.

Using matrix notation,

$$f(\mathbf{s}, \mathbf{x}) = \left[\frac{(1 - |\mathbf{x}|^2)(\pi_i(\mathbf{x}) - \mathbf{x}) \cosh^{-1} \frac{1 - \mathbf{x} \cdot \pi_i(\mathbf{x})}{\sqrt{1 - |\mathbf{x}|^2} \sqrt{1 - |\pi_i(\mathbf{x})|^2}}}{\sqrt{(\pi_i(\mathbf{x}) - \mathbf{x})^T ((1 - |\mathbf{x}|^2)I_n + \mathbf{x}\mathbf{x}^T)(\pi_i(\mathbf{x}) - \mathbf{x})}} \right]_{1 \leq i \leq n+1} \begin{bmatrix} \mathbf{s} \\ s_{n+1} \end{bmatrix},$$

which we abbreviate as $f(\mathbf{s}, \mathbf{x}) = [\mathbf{u}_1(\mathbf{x}) \ \dots \ \mathbf{u}_{n+1}(\mathbf{x})] \begin{bmatrix} \mathbf{s} \\ s_{n+1} \end{bmatrix}$.

We note again that although this function is undefined whenever \mathbf{x} lies in one of the flats, the singularities are removable.

Lemma 4.7. *Let Π_1, \dots, Π_{n+1} be pairwise disjoint flats in D^n . If $\mathbf{0} \notin \Pi_i$, then at $\mathbf{x} = \mathbf{0}$, the matrix $D\mathbf{u}_i$ is negative definite. If $\mathbf{0} \in \Pi_i$, then at $\mathbf{x} = \mathbf{0}$, the matrix $D\mathbf{u}_i$ is nonpositive definite. As a consequence, at $\mathbf{x} = \mathbf{0}$, if for all $i \leq n + 1$ we have that $0 \leq s_i < 1$, then the matrix*

$$\frac{\partial f}{\partial \mathbf{x}} = \begin{bmatrix} \frac{\partial f}{\partial x_1} & \dots & \frac{\partial f}{\partial x_n} \end{bmatrix}$$

is negative definite.

Proof. Assume first that $\mathbf{0} \notin \Pi_i$. To compute $D\mathbf{u}_i$ at $\mathbf{x} = \mathbf{0}$, we note that order-two terms in \mathbf{x} are irrelevant. Thus, we simplify $\mathbf{u}_i(\mathbf{x})$ to

$$\mathbf{u}_i(\mathbf{x}) = \frac{\pi_i(\mathbf{x}) - \mathbf{x}}{|\pi_i(\mathbf{x}) - \mathbf{x}|} \cosh^{-1} \frac{1 - \mathbf{x} \cdot \pi_i(\mathbf{x})}{\sqrt{1 - |\pi_i(\mathbf{x})|^2}} + \mathcal{O}(|\mathbf{x}|^2).$$

Recalling the expression for $\pi_i(\mathbf{x})$ from Proposition 2.3, we see that there are still more order-two terms within $|\pi_i(\mathbf{x})|^2$ and $\mathbf{x} \cdot \pi_i(\mathbf{x})$. Eliminating those, we simplify to

$$\mathbf{u}_i(\mathbf{x}) = \frac{\pi_i(\mathbf{x}) - \mathbf{x}}{|\pi_i(\mathbf{x}) - \mathbf{x}|} \cosh^{-1} \frac{1 - \mathbf{x} \cdot \pi_i(\mathbf{0})}{\sqrt{1 - |\pi_i(\mathbf{0})|^2}} + \mathcal{O}(|\mathbf{x}|^2).$$

The expression is now simple enough that we may compute $D\mathbf{u}_i$.

$$\begin{aligned} D\mathbf{u}_i &= \frac{D\pi_i(\mathbf{x}) - I_n}{|\pi_i(\mathbf{x}) - \mathbf{x}|} \cosh^{-1} \frac{1 - \mathbf{x} \cdot \pi_i(\mathbf{0})}{\sqrt{1 - |\pi_i(\mathbf{0})|^2}} \\ &\quad - \frac{\pi_i(\mathbf{x}) - \mathbf{x}}{|\pi_i(\mathbf{x}) - \mathbf{x}|^3} \left(\cosh^{-1} \frac{1 - \mathbf{x} \cdot \pi_i(\mathbf{0})}{\sqrt{1 - |\pi_i(\mathbf{0})|^2}} \right) (\pi_i(\mathbf{x}) - \mathbf{x})^T (D\pi_i(\mathbf{x}) - I_n) \\ &\quad + \frac{\pi_i(\mathbf{x}) - \mathbf{x}}{|\pi_i(\mathbf{x}) - \mathbf{x}|} \frac{1}{\sqrt{\left(\frac{1 - \mathbf{x} \cdot \pi_i(\mathbf{0})}{\sqrt{1 - |\pi_i(\mathbf{0})|^2}} \right)^2 - 1}} \frac{-\pi_i(\mathbf{0})^T}{\sqrt{1 - |\pi_i(\mathbf{0})|^2}} + \mathcal{O}(|\mathbf{x}|). \end{aligned}$$

Evaluating at $\mathbf{x} = \mathbf{0}$ gives

$$\begin{aligned} D\mathbf{u}_i|_{\mathbf{x}=\mathbf{0}} &= \frac{D\pi_i(\mathbf{0}) - I_n}{|\pi_i(\mathbf{0})|} \cosh^{-1} \frac{1}{\sqrt{1 - |\pi_i(\mathbf{0})|^2}} \\ &\quad - \frac{\pi_i(\mathbf{0})}{|\pi_i(\mathbf{0})|^3} \left(\cosh^{-1} \frac{1}{\sqrt{1 - |\pi_i(\mathbf{0})|^2}} \right) \pi_i(\mathbf{0})^T (D\pi_i(\mathbf{0}) - I_n) \\ &\quad - \frac{\pi_i(\mathbf{0})}{|\pi_i(\mathbf{0})|^2} \pi_i(\mathbf{0})^T \\ &= \frac{\tanh^{-1} |\pi_i(\mathbf{0})|}{|\pi_i(\mathbf{0})|} \left(I_n - \frac{\pi_i(\mathbf{0})\pi_i(\mathbf{0})^T}{|\pi_i(\mathbf{0})|^2} \right) (D\pi_i(\mathbf{0}) - I_n) - \frac{\pi_i(\mathbf{0})\pi_i(\mathbf{0})^T}{|\pi_i(\mathbf{0})|^2}. \end{aligned}$$

We must prove that for all i , $D\mathbf{u}_i|_{\mathbf{x}=\mathbf{0}}$ is negative definite. To do this, we decompose \mathbb{R}^n into a sum of three orthogonal vector spaces, $\mathbb{R}^n = \text{Span}(\pi_i(\mathbf{0})) \oplus T \oplus W$, where T is the k -dimensional tangent space to Π_i and W is the $(n - k - 1)$ -dimensional orthogonal complement to $\text{Span}(\pi_i(\mathbf{0})) \oplus T$. Then any vector $\mathbf{y} \in \mathbb{R}^n$ can be written in a unique way as $\mathbf{y} = \beta\pi_i(\mathbf{0}) + \mathbf{t} + \mathbf{w}$ for $\beta \in \mathbb{R}$, $\mathbf{t} \in T$, and $\mathbf{w} \in W$. By Lemma 2.5, $D\pi_i(\mathbf{0})\mathbf{y} = \alpha\mathbf{t}$ for some $0 < \alpha < 1$. Then

$$\begin{aligned} D\mathbf{u}_i|_{\mathbf{x}=\mathbf{0}}\mathbf{y} &= \frac{\tanh^{-1} |\pi_i(\mathbf{0})|}{|\pi_i(\mathbf{0})|} \left(I_n - \frac{\pi_i(\mathbf{0})\pi_i(\mathbf{0})^T}{|\pi_i(\mathbf{0})|^2} \right) (\alpha\mathbf{t} - \mathbf{y}) - \pi_i(\mathbf{0})\beta \\ &= \frac{\tanh^{-1} |\pi_i(\mathbf{0})|}{|\pi_i(\mathbf{0})|} (\alpha\mathbf{t} - \mathbf{y} + \pi_i(\mathbf{0})\beta) - \pi_i(\mathbf{0})\beta \\ &= \frac{\tanh^{-1} |\pi_i(\mathbf{0})|}{|\pi_i(\mathbf{0})|} ((\alpha - 1)\mathbf{t} - \mathbf{w}) - \pi_i(\mathbf{0})\beta. \end{aligned}$$

Thus,

$$\mathbf{y}^T D\mathbf{u}_i|_{\mathbf{x}=\mathbf{0}}\mathbf{y} = -|\pi_i(\mathbf{0})|^2\beta^2 - \frac{\tanh^{-1} |\pi_i(\mathbf{0})|}{|\pi_i(\mathbf{0})|} ((1 - \alpha)|\mathbf{t}|^2 + |\mathbf{w}|^2).$$

If $\mathbf{y} \neq \mathbf{0}$, then $\mathbf{y}^T D\mathbf{u}_i|_{\mathbf{x}=\mathbf{0}}\mathbf{y} < 0$.

Now assume that $\mathbf{0} \in \Pi_i$. We have already seen $\mathbf{u}_i(\mathbf{x})$ has a removable singularity at $\mathbf{x} = \mathbf{0}$, but now we must compute the derivative at $\mathbf{x} = \mathbf{0}$. Referring back to the form of $\pi_i(\mathbf{x})$ from Proposition 2.3, we see that $\mathbf{x} \cdot \pi_i(\mathbf{x}) = |\pi_i(\mathbf{x})|^2$. That allows us to simplify

$$\frac{(1 - |\mathbf{x}|^2)(\pi_i(\mathbf{x}) - \mathbf{x}) \cosh^{-1} \frac{1 - \mathbf{x} \cdot \pi_i(\mathbf{x})}{\sqrt{1 - |\mathbf{x}|^2}\sqrt{1 - |\pi_i(\mathbf{x})|^2}}}{\sqrt{(\pi_i(\mathbf{x}) - \mathbf{x})^T ((1 - |\mathbf{x}|^2)I_n + \mathbf{x}\mathbf{x}^T)(\pi_i(\mathbf{x}) - \mathbf{x})}}$$

to

$$\begin{aligned}
 \mathbf{u}_i(\mathbf{x}) &= \frac{(1 - |\mathbf{x}|^2)(\pi_i(\mathbf{x}) - \mathbf{x}) \cosh^{-1} \frac{1 - |\pi_i(\mathbf{x})|^2}{\sqrt{1 - |\mathbf{x}|^2}\sqrt{1 - |\pi_i(\mathbf{x})|^2}}}{\sqrt{(\pi_i(\mathbf{x}) - \mathbf{x})^T ((1 - |\mathbf{x}|^2)(\pi_i(\mathbf{x}) - \mathbf{x}) + \mathbf{x}(|\pi_i(\mathbf{x})|^2 - |\mathbf{x}|^2))}} \\
 &= \frac{(1 - |\mathbf{x}|^2)(\pi_i(\mathbf{x}) - \mathbf{x}) \cosh^{-1} \sqrt{\frac{1 - |\pi_i(\mathbf{x})|^2}{1 - |\mathbf{x}|^2}}}{\sqrt{(1 - |\mathbf{x}|^2)(|\pi_i(\mathbf{x})|^2 - 2|\pi_i(\mathbf{x})|^2 + |\mathbf{x}|^2) + (|\pi_i(\mathbf{x})|^2 - |\mathbf{x}|^2)^2}} \\
 &= \frac{(1 - |\mathbf{x}|^2)(\pi_i(\mathbf{x}) - \mathbf{x}) \tanh^{-1} \sqrt{\frac{|\mathbf{x}|^2 - |\pi_i(\mathbf{x})|^2}{1 - |\pi_i(\mathbf{x})|^2}}}{\sqrt{(|\mathbf{x}|^2 - |\pi_i(\mathbf{x})|^2)((1 - |\mathbf{x}|^2) + (|\mathbf{x}|^2 - |\pi_i(\mathbf{x})|^2))}} \\
 &= \frac{(1 - |\mathbf{x}|^2)(\pi_i(\mathbf{x}) - \mathbf{x})}{\sqrt{(|\mathbf{x}|^2 - |\pi_i(\mathbf{x})|^2)(1 - |\pi_i(\mathbf{x})|^2)}} \\
 &\quad \times \left(\sqrt{\frac{|\mathbf{x}|^2 - |\pi_i(\mathbf{x})|^2}{1 - |\pi_i(\mathbf{x})|^2}} + \mathbb{O}\left(\left(\frac{|\mathbf{x}|^2 - |\pi_i(\mathbf{x})|^2}{1 - |\pi_i(\mathbf{x})|^2}\right)^{3/2}\right) \right) \\
 &= \frac{(1 - |\mathbf{x}|^2)(\pi_i(\mathbf{x}) - \mathbf{x})}{1 - |\pi_i(\mathbf{x})|^2} \left(1 + \mathbb{O}\left(\frac{|\mathbf{x}|^2 - |\pi_i(\mathbf{x})|^2}{1 - |\pi_i(\mathbf{x})|^2}\right) \right).
 \end{aligned}$$

In computing the derivative at $\mathbf{x} = \mathbf{0}$, we may again ignore order-two terms in \mathbf{x} , meaning that the derivative at $\mathbf{x} = \mathbf{0}$ is simply $D\pi_i|_{\mathbf{x}=\mathbf{0}} - I_n$. Proceeding as before, this matrix is nonpositive definite on \mathbb{R}^n .

The \mathbf{x} -derivative of f is

$$\frac{\partial f}{\partial \mathbf{x}}|_{\mathbf{x}=\mathbf{0}} = \sum_{i=1}^{n+1} s_i D\mathbf{u}_i|_{\mathbf{x}=\mathbf{0}}.$$

The s_i add to 1, but they are all less than 1, so at least two of the s_i are positive. Since all of the matrices $D\mathbf{u}_i|_{\mathbf{x}=\mathbf{0}}$ are nonpositive definite, and at most one isn't negative definite, $\partial f/\partial \mathbf{x}|_{\mathbf{x}=\mathbf{0}}$ is negative definite. \square

Lemma 4.8. *Let Π_1, \dots, Π_{n+1} be pairwise disjoint flats in D^n . Suppose that $\mathbf{0}$ lies in their closed ideal Delaunay cell. If $\pi_1(\mathbf{0}), \dots, \pi_{n+1}(\mathbf{0})$ are in general position, then the matrix $\partial f/\partial \mathbf{s}|_{\mathbf{x}=\mathbf{0}} = [\partial f/\partial s_1 \ \dots \ \partial f/\partial s_n]|_{\mathbf{x}=\mathbf{0}}$ is invertible. The sign of the determinant of $\partial f/\partial \mathbf{s}|_{\mathbf{x}=\mathbf{0}}$ is $(-1)^n$ times the sign of the orientation of the simplex with vertices $\pi_1(\mathbf{0}), \dots, \pi_{n+1}(\mathbf{0})$.*

Proof. With $s_{n+1} = 1 - \sum_{i=1}^n s_i$, $\partial f/\partial s_i$ will be the coefficient of s_i minus the coefficient of s_{n+1} . At $\mathbf{x} = \mathbf{0}$, $\partial f/\partial s_i$ evaluates to

$$\begin{aligned} \frac{\partial f}{\partial s_i} \Big|_{\mathbf{x}=\mathbf{0}} &= \frac{\pi_i(\mathbf{0}) \cosh^{-1} \frac{1}{\sqrt{1-|\pi_i(\mathbf{0})|^2}}}{|\pi_i(\mathbf{0})|} - \frac{\pi_{n+1}(\mathbf{0}) \cosh^{-1} \frac{1}{\sqrt{1-|\pi_{n+1}(\mathbf{0})|^2}}}{|\pi_{n+1}(\mathbf{0})|} \\ &= \pi_i(\mathbf{0}) \frac{\tanh^{-1} |\pi_i(\mathbf{0})|}{|\pi_i(\mathbf{0})|} - \pi_{n+1}(\mathbf{0}) \frac{\tanh^{-1} |\pi_{n+1}(\mathbf{0})|}{|\pi_{n+1}(\mathbf{0})|}. \end{aligned}$$

Then

$$\text{Det} \frac{\partial f}{\partial \mathbf{s}} \Big|_{\mathbf{x}=\mathbf{0}} = * \bigwedge_{j=1}^n \left(\pi_j(\mathbf{0}) \frac{\tanh^{-1} |\pi_j(\mathbf{0})|}{|\pi_j(\mathbf{0})|} - \pi_{n+1}(\mathbf{0}) \frac{\tanh^{-1} |\pi_{n+1}(\mathbf{0})|}{|\pi_{n+1}(\mathbf{0})|} \right).$$

Since $\pi_1(\mathbf{0}), \dots, \pi_{n+1}(\mathbf{0})$ are in general position and $\mathbf{0}$ is in their convex hull, $\mathbf{0}$ has unique nonnegative Euclidean barycentric coordinates (t_1, \dots, t_{n+1}) and unique nonnegative hyperbolic barycentric coordinates (t'_1, \dots, t'_{n+1}) . These satisfy $t_i \bigwedge_{j=i+1}^{i+n} (\pi_j(\mathbf{0}) - \pi_i(\mathbf{0})) = \bigwedge_{j=i+1}^{i+n} \pi_j(\mathbf{0})$ and

$$t'_i \bigwedge_{j=i+1}^{i+n} \left(\pi_j(\mathbf{0}) \frac{\tanh^{-1} |\pi_j(\mathbf{0})|}{|\pi_j(\mathbf{0})|} - \pi_i(\mathbf{0}) \frac{\tanh^{-1} |\pi_i(\mathbf{0})|}{|\pi_i(\mathbf{0})|} \right) = \bigwedge_{j=i+1}^{i+n} \pi_j(\mathbf{0}) \frac{\tanh^{-1} |\pi_j(\mathbf{0})|}{|\pi_j(\mathbf{0})|}.$$

Obviously, $\bigwedge_{j=i+1}^{i+n} \pi_j(\mathbf{0})$ and $\bigwedge_{j=i+1}^{i+n} \pi_j(\mathbf{0}) (\tanh^{-1} |\pi_j(\mathbf{0})|) / |\pi_j(\mathbf{0})|$ are positive scalar multiples of each other. There is some i for which t_i is positive. Since $\bigwedge_{j=i+1}^{i+n} (\pi_j(\mathbf{0}) - \pi_i(\mathbf{0}))$ is nonzero,

$$t'_i \bigwedge_{j=i+1}^{i+n} \left(\pi_j(\mathbf{0}) \frac{\tanh^{-1} |\pi_j(\mathbf{0})|}{|\pi_j(\mathbf{0})|} - \pi_i(\mathbf{0}) \frac{\tanh^{-1} |\pi_i(\mathbf{0})|}{|\pi_i(\mathbf{0})|} \right)$$

must also be nonzero. Then, for this i , $\bigwedge_{j=i+1}^{i+n} (\pi_j(\mathbf{0}) - \pi_i(\mathbf{0}))$ and

$$\bigwedge_{j=i+1}^{i+n} \left(\pi_j(\mathbf{0}) \frac{\tanh^{-1} |\pi_j(\mathbf{0})|}{|\pi_j(\mathbf{0})|} - \pi_i(\mathbf{0}) \frac{\tanh^{-1} |\pi_i(\mathbf{0})|}{|\pi_i(\mathbf{0})|} \right)$$

are positive scalar multiples of each other. By Lemma 3.2, this is true for all i . The sign of $\text{Det} \partial f / \partial \mathbf{s} |_{\mathbf{x}=\mathbf{0}}$ is the same as the sign of $* \bigwedge_{j=1}^n (\pi_j(\mathbf{0}) - \pi_{n+1}(\mathbf{0}))$, and so is the orientation of the simplex with vertices $\pi_{n+1}(\mathbf{0}), \pi_1(\mathbf{0}), \dots, \pi_n(\mathbf{0})$. Reordering the vertices to $\pi_1(\mathbf{0}), \dots, \pi_{n+1}(\mathbf{0})$ changes the sign by a factor of $(-1)^n$. \square

Definition 4.9. The standard n -dimensional simplex is the subset of \mathbb{R}^{n+1} given by

$$\left\{ (s_1, \dots, s_{n+1}) : \sum_{i=1}^{n+1} s_i = 1 \text{ and for all } i, s_i \geq 0 \right\}.$$

The standard ideal n -dimensional simplex Δ_n is the standard n -dimensional simplex with its vertices removed. For purposes of orientation, the order of the vertices is $(1, 0, \dots, 0), (0, 1, 0, \dots, 0), \dots, (0, \dots, 0, 1)$.

Proposition 4.10. *Given pairwise disjoint flats Π_1, \dots, Π_{n+1} in D^n , let X be the closed ideal Delaunay cell associated with these flats. Let $\overset{\circ}{X}$ be the open Delaunay cell associated with these flats. There is a surjective differentiable map $\phi : \Delta_n \rightarrow X$ such that*

- (1) $\phi : \phi^{-1}(\overset{\circ}{X}) \rightarrow \overset{\circ}{X}$ is a diffeomorphism (assuming $\overset{\circ}{X}$ is nonempty),
- (2) for $\mathbf{t} \in \phi^{-1}(\overset{\circ}{X})$, ϕ is orientation-preserving if and only if the simplex with vertices $\pi_1(\phi(\mathbf{t})), \dots, \pi_{n+1}(\phi(\mathbf{t}))$ has positive orientation, and
- (3) ϕ carries m -dimensional faces of Δ_n to m -dimensional closed ideal Delaunay cells within X . Restricted to any m -dimensional face of Δ_n , ϕ depends on only $m + 1$ of the $n + 1$ flats.

Proof. By Lemma 4.7, we know that if $0 \leq s_i < 1$ for $i \leq n + 1$, then $\partial f / \partial \mathbf{x}$ is invertible at $\mathbf{x} = \mathbf{0}$. Hyperbolic barycentric coordinates are preserved by hyperbolic isometries, and the function f is the defining equation for computing hyperbolic barycentric coordinates, so without loss of generality, we may perform a hyperbolic isometry moving any point of interest to the origin. Thus $\partial f / \partial \mathbf{x}$ is invertible at every point in D^n .

A point $\mathbf{x} \in D^n$ is in X if and only if it lies in some open m -dimensional Delaunay cell. This is equivalent to \mathbf{x} having positive hyperbolic barycentric coordinates with respect to some nonsingleton subset of $\pi_1(\mathbf{x}), \dots, \pi_{n+1}(\mathbf{x})$ and $\mathbf{0}$ for the remaining coordinates. Thus, \mathbf{x} is in X if and only if there is some $\mathbf{t} \in \Delta_n$ with

$$\mathbf{t} = \begin{bmatrix} \mathbf{s} \\ s_{n+1} \end{bmatrix}$$

for which $f(\mathbf{s}, \mathbf{x}) = \mathbf{0}$.

We claim that there is a unique $\mathbf{x} \in X$ such that $f((\frac{1}{2}, \frac{1}{2}, 0, \dots, 0), \mathbf{x}) = \mathbf{0}$. At any point $\mathbf{x} \in X$ for which $f((\frac{1}{2}, \frac{1}{2}, 0, \dots, 0), \mathbf{x}) = \mathbf{0}$, the vectors (in $T_x \mathbb{H}^n$) to $\pi_1(\mathbf{x})$ and $\pi_2(\mathbf{x})$ must be equal and opposite. Thus, \mathbf{x} must be the midpoint of the common perpendicular to Π_1 and Π_2 .

Thus, we meet the hypotheses for Lemma 4.5, with

$$A = \{(s_1, \dots, s_n) \in \mathbb{R}^n : 0 \leq s_i < 1 \text{ for } i \leq n + 1\}$$

and $B = X$. Then there is a surjective differentiable function $g : A \rightarrow X$ such that $f(\mathbf{s}, \mathbf{x}) = \mathbf{0}$ if and only if $\mathbf{x} = g(\mathbf{s})$.

It is a local diffeomorphism at any point where $\partial f/\partial s$ has rank n . By Lemma 4.8, g is a local diffeomorphism at each point s for which the affine hull of

$$\pi_1(g(s)), \dots, \pi_{n+1}(g(s))$$

is n -dimensional. Thus, g is a local diffeomorphism at each point in $g^{-1}(\overset{\circ}{X})$. Further, each point in $\overset{\circ}{X}$ has unique hyperbolic barycentric coordinates, proving that $g : g^{-1}(\overset{\circ}{X}) \rightarrow \overset{\circ}{X}$ is a bijection.

Let $\phi : \Delta_n \rightarrow A \rightarrow X$ be the composite of g and projection onto the first n coordinates. Such a projection is a diffeomorphism, so we have proved the first claim.

The derivative of g is given by $\partial g/\partial \mathbf{x} = (-\partial f/\partial \mathbf{x})^{-1} \partial f/\partial s$. Because $-\partial f/\partial \mathbf{x}$ is positive definite, it has positive determinant. The determinant of $\partial f/\partial s$ has sign equal to $(-1)^n$ times the sign of the orientation of the simplex with vertices $\pi_1(g(s)), \dots, \pi_{n+1}(g(s))$. Thus, $\text{Det } \partial g/\partial \mathbf{x}$ has sign equal to $(-1)^n$ times the sign of the orientation of the simplex with vertices $\pi_1(g(s)), \dots, \pi_{n+1}(g(s))$. Since the projection map from Δ_n to \mathbb{R}^n is orientation-preserving/reversing when n is even/odd respectively, ϕ is orientation-preserving/reversing depending on the sign of the orientation of the simplex with vertices $\pi_1(\phi(t)), \dots, \pi_{n+1}(\phi(t))$.

To prove the third claim, suppose without loss of generality that the first $m+1$ coordinates of t are the only nonzero coordinates. Then $\phi(t)$ lies in the convex hull of $\pi_1(\mathbf{0}), \dots, \pi_{m+1}(\mathbf{0})$, and so it lies within a closed m -dimensional Delaunay cell. Further, the $i = m+2$ through $i = n+1$ terms in f are all 0, so f does not depend on $\Pi_{m+2}, \dots, \Pi_{n+1}$, and thus ϕ doesn't either. \square

Definition 4.11. Let T_n be a truncated n -dimensional simplex. Let T'_n be T_n with the truncation faces removed. T'_n is diffeomorphic to Δ_n .

Theorem 4.12. *Given pairwise disjoint flats Π_1, \dots, Π_{n+1} in D^n , let X be the closed ideal Delaunay cell associated with these flats. There is a surjective continuous function $\bar{\phi} : T_n \rightarrow \bar{X}$ such that when restricted to T'_n , $\bar{\phi}$ is equivalent to the function ϕ given by Proposition 4.10. The truncation faces are mapped to subsets of the flats Π_1, \dots, Π_{n+1} . For any particular truncation face, the map $\bar{\phi}$ can be chosen so the level sets along that truncation face are either affine subspaces of dimension at least $n-k-1$ or lie in the convex hull of some of the truncation face's vertices. For any particular truncation face, if $k=1$, the map $\bar{\phi}$ can be chosen so the level sets along that truncation face either are affine subspaces of dimension at least $n-2$ or equal the convex hull of some of the truncation face's vertices.*

Proof. Without loss of generality, we may perform a hyperbolic isometry so that $\mathbf{0} \in \Pi_{n+1}$. Then $\pi_{n+1} : D^n \rightarrow \Pi_{n+1}$ is just the Euclidean projection from D^n to Π_{n+1} . In particular, it's a linear transformation. Thus, $D\pi_{n+1}$ is a constant

projection matrix and $\pi_{n+1}(\mathbf{x}) = D\pi_{n+1}\mathbf{x}$ for any $\mathbf{x} \in D^n$. From this, we see that $(D\pi_{n+1})(\pi_{n+1}(\mathbf{x}) - \mathbf{x}) = \mathbf{0}$.

As in the preceding proof, we may project Δ_n to $A \subset \mathbb{R}^n$. Let $h : A \rightarrow \Delta_n$ be the inverse of this projection. Then $g = \phi \circ h$. Consider the map $r : \{\mathbf{v} \in \mathbb{R}^n : |\mathbf{v}| \geq 1\} \rightarrow \mathbb{R}^n$ given by $r(\mathbf{v}) = (1 - 1/|\mathbf{v}|)\mathbf{v}$. This map moves each point in $\{\mathbf{v} \in \mathbb{R}^n : |\mathbf{v}| \geq 1\}$ one unit closer to the origin. The closure of $r^{-1}(A)$ is diffeomorphic to an n -dimensional simplex with one vertex truncated along a face F . The face F is the face of the closure of $r^{-1}(A)$ that is nearest the origin, so it is the portion of the unit sphere that has all of its coordinates nonnegative.

We construct a function $f' : F \times \Pi_{n+1} \rightarrow \Pi_{n+1}$ with

$$f'(\mathbf{v}, \mathbf{x}) = D\pi_{n+1}[\mathbf{u}_1(\mathbf{x}) \dots \mathbf{u}_n(\mathbf{x})]\mathbf{v}.$$

Let $B \subset \Pi_{n+1}$ be the set

$$B = \{\mathbf{x} \in \Pi_{n+1} : \text{there is some } \mathbf{v} \in F \text{ for which } f'(\mathbf{v}, \mathbf{x}) = \mathbf{0}\}.$$

Letting $\mathbf{v} = (v_1, \dots, v_n)$, the matrix $\partial f'/\partial \mathbf{x}$ is $\sum_{j=1}^n v_j(D\pi_{n+1})(D\mathbf{u}_j)$. With \mathbf{u}_j viewed as a function defined on D^n , $D\mathbf{u}_j$ is negative definite at any $\mathbf{x} \in \Pi_{n+1}$, so when restricted to the tangent space to Π_{n+1} , $(D\pi_{n+1})(D\mathbf{u}_j)$ is also negative definite, and so has rank k . Thus, $\partial f'/\partial \mathbf{x}$ is invertible. If $f'((1, 0, \dots, 0), \mathbf{x}) = \mathbf{0}$, then $(D\pi_{n+1})\mathbf{u}_1(\mathbf{x}) = \mathbf{0}$, so at \mathbf{x} , the geodesic to the closest point on Π_1 is perpendicular to Π_{n+1} . Of course, that same geodesic must be perpendicular to Π_1 , so it must be their common perpendicular. Thus, there is a unique $\mathbf{x} \in \Pi_{n+1}$ such that $f'((1, 0, \dots, 0), \mathbf{x}) = \mathbf{0}$. Thus, we meet the hypotheses for Lemma 4.5, so there is a surjective differentiable function $\bar{\phi}' : F \rightarrow B \subset \Pi_{n+1}$ such that $f'(\mathbf{v}, \mathbf{x}) = \mathbf{0}$ if and only if $\mathbf{x} = \bar{\phi}'(\mathbf{v})$. For any point in $\mathbf{x} \in B$, the preimage under $\bar{\phi}'$ is the set of all $\mathbf{v} \in F$ such that $f'(\mathbf{v}, \mathbf{x}) = \mathbf{0}$. The matrix $D\pi_{n+1}[\mathbf{u}_1(\mathbf{x}) \dots \mathbf{u}_n(\mathbf{x})]$ is a $k \times n$ matrix, so its nullspace (in \mathbb{R}^n) has dimension at least $n - k$. The intersection of the nullspace with the first orthant either has the same dimension as the nullspace or lies entirely within the span of some of the coordinate axes. The preimage of \mathbf{x} through $\bar{\phi}'$ is the intersection of the nullspace with the portion of the unit sphere in the first orthant, and so it either has dimension at least $n - k - 1$ or lies in the (spherical) convex hull of some of the vertices of F . In the case that $k = 1$, the nullspace has dimension at least $n - 1$, so its intersection with the first orthant either is of dimension at least $n - 1$ or equals the span of some of the coordinate axes.

Define a function $\bar{\phi} : r^{-1}(A) \cup F \rightarrow X \cup \Pi_{n+1}$ by

$$\bar{\phi}(\mathbf{v}) = \begin{cases} \bar{\phi}'(\mathbf{v}) & \text{if } \mathbf{v} \in F, \\ g(r(\mathbf{v})) & \text{if } \mathbf{v} \in r^{-1}(A). \end{cases}$$

Consider any sequence $\mathbf{v}_j \in r^{-1}(A)$ that converges to a point on $\mathbf{v} \in F$. Then $\lim_{j \rightarrow \infty} r(\mathbf{v}_j) = \mathbf{0}$. Since $\bar{\phi}(\mathbf{v}_j) \in X \subset D^n$, $\bar{\phi}(\mathbf{v}_j)$ has a convergent subsequence that

converges to a point $\mathbf{x} \in D^n$. Without loss of generality, assume $\lim_{j \rightarrow \infty} \bar{\phi}(\mathbf{v}_j) = \mathbf{x}$. We claim that $\mathbf{x} = \bar{\phi}(\mathbf{v})$.

Based on the definition of ϕ , $f(r(\mathbf{v}_j), g(r(\mathbf{v}_j))) = \mathbf{0}$. Taking a limit, $f(\mathbf{0}, \mathbf{x}) = \mathbf{0}$. Thus, $\pi_{n+1}(\mathbf{x}) = \mathbf{x}$, so \mathbf{x} must lie in Π_{n+1} .

It must also be true that

$$\frac{1}{|\mathbf{v}_j| - 1} (D\pi_{n+1})f(r(\mathbf{v}_j), g(r(\mathbf{v}_j))) = \frac{1}{|\mathbf{v}_j| - 1} (D\pi_{n+1})\mathbf{0} = \mathbf{0}.$$

Using the matrix version of f ,

$$f(r(\mathbf{v}_j), g(r(\mathbf{v}_j))) = [\mathbf{u}_1(g(r(\mathbf{v}_j))) \ \dots \ \mathbf{u}_{n+1}(g(r(\mathbf{v}_j)))]h(r(\mathbf{v}_j)).$$

The vector $\mathbf{u}_{n+1}(g(r(\mathbf{v}_j)))$ is a scalar multiple of $\pi_{n+1}(g(r(\mathbf{v}_j))) - g(r(\mathbf{v}_j))$, so $(D\pi_{n+1})\mathbf{u}_{n+1}(g(r(\mathbf{v}_j))) = \mathbf{0}$. Thus, when computing $(D\pi_{n+1})f(r(\mathbf{v}_j), g(r(\mathbf{v}_j)))$, we may ignore the final column of $[\mathbf{u}_1(g(r(\mathbf{v}_j))) \ \dots \ \mathbf{u}_{n+1}(g(r(\mathbf{v}_j)))]$ and the final entry of the vector the matrix is multiplied by, $h(r(\mathbf{v}_j))$. Ignoring the final entry of $h(r(\mathbf{v}_j))$ produces $r(\mathbf{v}_j)$. This simplifies

$$\mathbf{0} = \frac{1}{|\mathbf{v}_j| - 1} (D\pi_{n+1})f(r(\mathbf{v}_j), g(r(\mathbf{v}_j)))$$

to

$$\begin{aligned} \mathbf{0} &= \frac{1}{|\mathbf{v}_j| - 1} D\pi_{n+1}[\mathbf{u}_1(g(r(\mathbf{v}_j))) \ \dots \ \mathbf{u}_n(g(r(\mathbf{v}_j)))]r(\mathbf{v}_j) \\ &= \frac{1}{|\mathbf{v}_j| - 1} D\pi_{n+1}[\mathbf{u}_1(g(r(\mathbf{v}_j))) \ \dots \ \mathbf{u}_n(g(r(\mathbf{v}_j)))]\left(1 - \frac{1}{|\mathbf{v}_j|}\right)\mathbf{v}_j \\ &= D\pi_{n+1}[\mathbf{u}_1(g(r(\mathbf{v}_j))) \ \dots \ \mathbf{u}_n(g(r(\mathbf{v}_j)))]\frac{\mathbf{v}_j}{|\mathbf{v}_j|}. \end{aligned}$$

Taking a limit as $j \rightarrow \infty$, we have $\mathbf{0} = D\pi_{n+1}[\mathbf{u}_1(\mathbf{x}) \ \dots \ \mathbf{u}_n(\mathbf{x})]\mathbf{v}$, so

$$f'(\mathbf{v}, \mathbf{x}) = \mathbf{0}.$$

Thus, $\mathbf{x} = \bar{\phi}'(\mathbf{v}) = \bar{\phi}(\mathbf{v})$. We have proved that $\bar{\phi}$ is continuous along F . Following the same process, we may extend $\bar{\phi}$ to each truncation face in a continuous fashion. Since each point on a truncation face is a limit of points in T'_n , we have that $\bar{\phi}(T_n)$ is contained within the closure of $\bar{\phi}(T'_n) = X$. Since a truncated n -simplex is compact, its image under a continuous map is compact as well. Thus, $\bar{\phi}(T_n) = \bar{X}$. \square

5. The Voronoi decomposition

In order to construct a Delaunay decomposition, we first need to construct a suitable Voronoi decomposition. Although the Voronoi decomposition can be constructed for a wide variety of different kinds of arrangements of objects, we restrict attention to symmetric cocompact arrangements of k -dimensional flats in \mathbb{H}^n .

To construct a Delaunay complex, we use an abstraction of the Voronoi complex. The Voronoi complex, as usually defined, associates to each $\Pi \in \mathcal{P}$ a Voronoi cell, given by the set of points that are closer to Π than to any other flat in \mathcal{P} . Separating such cells from each other are Voronoi faces of various dimensions less than n . Each face is equidistant from some finite number of the flats in \mathcal{P} . Thus, a face could be labeled with a finite nonsingleton proper subset of \mathcal{P} . An abstract Voronoi complex has many of the properties of the traditional Voronoi complex, but has fewer geometric criteria and more topological criteria.

Definition 5.1. For a symmetric cocompact arrangement \mathcal{P} of k -dimensional flats in \mathbb{H}^n , an abstract simple Voronoi complex is a Γ -equivariant partition \mathcal{V} of \mathbb{H}^n into connected submanifolds (without boundary) labeled by finite nonempty subsets of \mathcal{P} such that:

- (1) The Γ -action on the labels is compatible with the Γ -action on \mathcal{P} .
- (2) Ignoring labels, there are only finitely many Γ -orbits in \mathcal{V} .
- (3) For each $F \in \mathcal{V}$, the set $\bar{F} - F$ is of dimension $\dim(F) - 1$. We call this set the boundary of F and denote it by ∂F .
- (4) For each j -dimensional $F \in \mathcal{V}$, the label for F has length at least $n + 1 - j$.
- (5) For each n -dimensional $F \in \mathcal{V}$, the pair $(F \cup \partial F, \partial F)$ is homeomorphic to the pair $(\mathbb{R}^k \times B^{n-k}, \mathbb{R}^k \times S^{n-k-1})$. Further, the label for F is $\{\Pi\} \subset \mathcal{P}$ for some $\Pi \subset \mathcal{P}$.
- (6) For each $F \in \mathcal{V}$ of dimension $j < n$, the pair $(F \cup \partial F, \partial F)$ is homeomorphic to the pair (B^j, S^{j-1}) . Such F are referred to as faces.
- (7) For $F, G \in \mathcal{V}$, if F intersects the boundary of G , then F is contained within the boundary of G . Further, the label for F properly contains the label for G .
- (8) Given $F \in \mathcal{V}$ and a nonempty proper subset A of the label of F , there is at most one $G \in \mathcal{V}$ such that $F \subset \partial G$ and G is labeled A . (There might be another $G' \in \mathcal{V}$ also labeled A , but it can't contain F in its boundary.)
- (9) Let $S \subset \mathbb{H}^n$ be the union of all faces of \mathcal{V} . Then $\mathbb{H}^n - \bigcup_{\Pi \in \mathcal{P}} \Pi$ is a regular neighborhood of S , and S is a deformation retract of $\mathbb{H}^n - \bigcup_{\Pi \in \mathcal{P}} \Pi$. The retraction map can be chosen to be Γ -equivariant. Thus, S/Γ is a deformation retract of $(\mathbb{H}^n - \bigcup_{\Pi \in \mathcal{P}} \Pi)/\Gamma$.

It's not clear that every symmetric cocompact arrangement \mathcal{P} has an abstract simple Voronoi complex. In general, we can't be sure that the faces of the Voronoi decomposition are submanifolds at all. The possibility exists that some of the faces might have singularities. Even if the faces are submanifolds, they might not be homeomorphic to disks.

However, there are some circumstances under which the Voronoi decomposition is an abstract simple Voronoi complex.

Proposition 5.2. *For a symmetric cocompact arrangement \mathcal{P} of k -dimensional flats in \mathbb{H}^n with either $k = n - 1$ or $n = 3, k = 1$, the Voronoi decomposition is an abstract simple Voronoi complex.*

Proof. Using somewhat different terminology, Dowty [2000] proved that for $n = 3$ and $k = 1$, the Voronoi decomposition satisfies all of the criteria for an abstract simple Voronoi complex except possibly criteria (7) and (8). In addition, he proved that each 2-dimensional face has a label of length exactly 2.

In the case $k = n - 1$, the set of points equidistant from two $(n - 1)$ -dimensional flats is itself an $(n - 1)$ -dimensional flat, from which one quickly deduces that for $j < n$, all j -dimensional faces of the Voronoi decomposition are the relative interiors of convex polytopes and that their labels have length at least $n + 1 - j$. Dowty's proofs of criteria (2), (5), and (9) easily generalize beyond $n = 3, k = 1$. His proof that each 2-dimensional face in \mathbb{H}^3 has a label of length 2 easily generalizes to prove that each $(n - 1)$ -dimensional face in \mathbb{H}^n has a label of length exactly 2. This leaves us with the criteria (7) and (8) to prove in both cases.

At any point \mathbf{x} on a j -dimensional face F labeled $\{\Pi_1, \dots, \Pi_m\}$, let the normal space to F be denoted $N_{\mathbf{x}}F$. Within $N_{\mathbf{x}}F$, consider the set of vectors along which a small perturbation moves into the cell labeled $\{\Pi_i\}$. This set is a convex radial $(n - j)$ -dimensional subset of $N_{\mathbf{x}}F$. Thus, allowing i to vary from 1 to m , we produce a tessellation of $N_{\mathbf{x}}F$ into m different convex radial regions. Each of these regions corresponds to one of the cells labeled $\{\Pi_i\}$, and meetings of two or more of these regions correspond to faces of \mathcal{V} labeled by subsets of $\{\Pi_1, \dots, \Pi_m\}$. Thus, if we can prove that this tessellation doesn't vary (in a combinatorial sense) with \mathbf{x} , we will have proved criteria (7) and (8).

For an $(n - 1)$ -dimensional face, this tessellation simply breaks the normal space into two rays, and so it is obviously independent of \mathbf{x} . For an $(n - 2)$ -dimensional face, the tessellation breaks the 2-dimensional normal space into m different radial sectors. Although the size of the sectors might vary with \mathbf{x} , the order of the sectors doesn't vary, since all m sectors must be present for every \mathbf{x} . For a 0-dimensional face (a vertex), there's nothing to prove, as there's only one point in the face. This completes the proof for the case $n = 3, k = 1$.

What remains to be considered is the case $k = n - 1$. Let F be a j -dimensional face of the Voronoi tessellation. Then F is an open subset of some j -dimensional flat. We may use the Klein model for \mathbb{H}^n and without loss of generality, we may assume that $\mathbf{0} \in F$. One of the basic properties of codimension-1 flats in the Klein model is that for any $\mathbf{x} \in D^n$, the line joining \mathbf{x} to $\pi_i(\mathbf{x}) \in \Pi_i$ must pass through the point $\pi_i(\mathbf{0})/|\pi_i(\mathbf{0})|^2$, so the direction vector from \mathbf{x} to $\pi_i(\mathbf{x})$ is a scalar multiple

of $\mathbf{x} - \pi_i(\mathbf{0})/|\pi_i(\mathbf{0})|^2$. Let \mathbf{x} be a point on F . Since we have chosen $\mathbf{0}$ to lie on F , the set $N_{\mathbf{x}}F$ is the same regardless of whether we view it in a Euclidean or hyperbolic sense. If we take the direction vector from \mathbf{x} to $\pi_i(\mathbf{x})$ and project it into $N_{\mathbf{x}}F$, we are left with a scalar multiple of the projection of $\pi_i(\mathbf{0})$ into $N_{\mathbf{x}}F$. This verifies that at every point $\mathbf{x} \in F$, the direction vectors from \mathbf{x} to $\pi_1(\mathbf{x}), \dots, \pi_m(\mathbf{x})$, when projected into $N_{\mathbf{x}}F$, are independent of \mathbf{x} , so the tessellation of $N_{\mathbf{x}}F$ is also independent of \mathbf{x} . □

In some specific cases, we have verified that the Voronoi decomposition produces an abstract simple Voronoi complex. However, even in those cases, we need stricter criteria to proceed. The Voronoi decomposition cannot be assumed to meet the new criteria, but as we will see, it can be modified to meet them.

Definition 5.3. We say that a face F in an abstract simple Voronoi complex is standard if

- (1) the cross-sectional tessellation used in the preceding proof is (combinatorially) the same at every point in F ,
- (2) each face containing F in its boundary is standard, and
- (3) for each nonempty proper subset A of the label of F , there is a unique $G \in \mathcal{V}$ such that $F \subset \partial G$ and G is labeled A .

The reason for using the word “standard” is that we want S/Γ to be a standard spine for $(\mathbb{H}^n - \bigcup_{\Pi \in \mathcal{P}} \Pi)/\Gamma$.

Proposition 5.4. *Let F be a standard face in an abstract simple Voronoi complex, and let the label for F be $\{\Pi_1, \dots, \Pi_m\}$. Then F is of dimension exactly $n + 1 - m$. The link of F is the boundary of an $(m - 1)$ -dimensional simplex whose vertices are labeled $\{\Pi_1\}, \dots, \{\Pi_m\}$ and whose faces are labeled by the union of their vertices’ labels.*

Proof. By the definition of an abstract simple Voronoi complex, the dimension of F is $j \geq n + 1 - m$. By the definition of a standard face, F lies in the boundary of a standard face labeled $\{\Pi_1, \dots, \Pi_{m-1}\}$, which itself lies in the boundary of a standard face labeled $\{\Pi_1, \dots, \Pi_{m-2}\}$, etc., which all lie in the boundary of the n -dimensional cell labeled $\{\Pi_1\}$. Since the boundary of a face must have lower dimension than the face itself, we see that F has dimension $j \leq n + 1 - m$. Thus, F has dimension exactly $j = n + 1 - m$.

The cross-sectional tessellation consists of $(m - 1)$ -dimensional radial sectors labeled by the various length-1 subsets of $\{\Pi_1, \dots, \Pi_m\}$, which meet along $(m - 2)$ -dimensional facets labeled by the various length-2 subsets of $\{\Pi_1, \dots, \Pi_m\}$, etc. Each face has the property that its label is the union of the labels of the cells containing that face in their boundary. The dual of this tessellation is a simplex consisting of vertices labeled by the various length-1 subsets of $\{\Pi_1, \dots, \Pi_m\}$,

connected by edges labeled by the union of their two vertices' labels, etc. The link is the boundary of that simplex. \square

Proposition 5.5. *Let F be a nonstandard face in an abstract simple Voronoi complex and let the label for F be $\{\Pi_1, \dots, \Pi_m\}$. If every face that contains F in its boundary is standard and the cross-sectional tessellation is (combinatorially) the same at every point in F , then the dimension of F is strictly greater than $n + 1 - m$.*

Proof. The definition of an abstract simple Voronoi complex requires that the dimension of F be at least $n + 1 - m$. Suppose that the dimension of F is exactly $n + 1 - m$. The cross-sections are of dimension $m - 1$ and are broken into m convex radial regions, one for each Π_i . Each j -dimensional cell of this tessellation is the cross-section of an $(n + 1 - m + j)$ -dimensional cell of \mathcal{V} . By Proposition 5.4, for $j > 0$, the length of the label of that face is $(n + 1) - (n + 1 - m + j) = m - j$. In particular, each edge of this tessellation has a label that is a length- $(m - 1)$ subset of $\{\Pi_1, \dots, \Pi_m\}$. By repeated application of criterion (3) in Definition 5.1, there are edges in the tessellation. Without loss of generality, there is an edge labeled $\{\Pi_2, \dots, \Pi_m\}$. However, by assumption, F is not standard, so for some proper subset $A \subset \{\Pi_1, \dots, \Pi_m\}$, there is no cell of the tessellation labeled A . Without loss of generality, $\Pi_m \notin A$. Then there is no edge of the tessellation labeled $\{\Pi_1, \dots, \Pi_{m-1}\}$, for if there were, the assumption that only standard faces contain F in their boundary would imply that there is a face labeled A that contains F in its boundary, and is thus a cell of the tessellation labeled A .

By the standardness assumption for all faces containing F in their boundary, the fact that there is an edge of the tessellation labeled $\{\Pi_2, \dots, \Pi_m\}$ means that there is a 2-dimensional cell of the tessellation labeled $\{\Pi_2, \dots, \Pi_{m-1}\}$. Since the cells of the tessellation are convex radial regions, that 2-dimensional cell must have a second edge on its boundary. The label for that edge must properly contain $\{\Pi_2, \dots, \Pi_{m-1}\}$ and must be different from $\{\Pi_2, \dots, \Pi_m\}$. Thus, there is an edge labeled $\{\Pi_1, \dots, \Pi_{m-1}\}$, which is a contradiction. \square


Definition 5.6. A standard abstract Voronoi complex is an abstract simple Voronoi complex in which every face is standard.

Proposition 5.7. *A symmetric cocompact arrangement \mathcal{P} of k -dimensional flats in \mathbb{H}^n with either $k = n - 1$ or $n = 3, k = 1$ has a standard abstract Voronoi complex.*


Proof. If the Voronoi decomposition is a standard abstract Voronoi complex, then we're already done. Otherwise, create an abstract simple Voronoi complex \mathcal{V} that is identical to the Voronoi decomposition. For a given face, the cross-sectional tessellation is (combinatorially) the same at every point on the face. There are only finitely many Γ -orbits of faces.

There is some nonstandard face. Among the Γ -orbits of nonstandard faces, pick a face F of the largest possible dimension. Let the label for F be $\{\Pi_1, \dots, \Pi_m\}$. Then F meets the criteria for Proposition 5.5, and so has dimension $j > n + 1 - m$.

At each $x \in F \cup \partial F$, take a small closed cross-section transverse to F in such a way that the union C (see figure on the right, with F as the central line) of these cross-sections contains a neighborhood of F , intersects no other cells of \mathcal{V} except those that either contain F within their boundary or are contained within the boundary of F , and is homeomorphic to $\overline{B^{n-j}} \times (F \cup \partial F)$.




By performing some homeomorphism within C , we may assume that the cross-sectional tessellation (illustrated below) is exactly the same for every $x \in F$.




Each cell of the tessellation is labeled by a proper subset of $\{\Pi_1, \dots, \Pi_m\}$. There is some proper subset of $\{\Pi_1, \dots, \Pi_m\}$ that is not the label of any of the cells of the tessellation. However, every singleton subset of $\{\Pi_1, \dots, \Pi_m\}$ is the label of exactly one $(n - j)$ -dimensional cell of the tessellation.

The dual within B^{n-j} of the tessellation is a polytope K . The boundary of K




(the link of F) is a triangulation of S^{n-j-1} , because every cell of \mathcal{V} of dimension greater than j is standard. The interior of K is labeled $\{\Pi_1, \dots, \Pi_m\}$, and the faces and vertices are labeled by some, but not all, of the nonempty proper subsets of $\{\Pi_1, \dots, \Pi_m\}$. Each face of K is labeled by the union of its vertices' labels. There are exactly m vertices of K and they are labeled $\{\Pi_1\}, \dots, \{\Pi_m\}$. Combinatorially triangulate K without adding any new vertices. This subdivides the interior of K into several cells of dimension $n - j$:




Label each of these new cells (of any dimension) with the union of the labels of its vertices. Thus, this triangulation has the property that every cell of any dimension is a simplex labeled by the union of the labels of its vertices. Thus, given any simplex, for each nonempty proper subset of its label, there is a unique simplex in its boundary bearing that subset as its label.

Finally, take the dual of this triangulation to produce a new tessellation of B^{n-j} :



Since the only changes we made to K were in the interior of K , we see that along S^{n-j-1} , this new tessellation is the same as the original tessellation. Within the interior of B^{n-j} , we have added some new cells of dimension less than $n - j$. Let C' be the product of this new tessellation with $F \cup \partial F$ (see figure on the right). Along $B^{n-j} \times \partial F$, perform a quotient map within each copy of B^{n-j} to identify all of the new cells to a single vertex. Let C'' be C' under this quotient map. Thus, C'' and C are identical along their boundaries. Within \mathcal{V} , replace C by C'' . The nonstandard face F has been replaced by various standard faces (not all of the same dimension). Using the Γ -action, make the same change at every face in the Γ -orbit of F . We now have an abstract simple Voronoi complex in which the number of Γ -orbits of nonstandard faces has been decreased by 1. Continue until all faces are standard. □



6. The Delaunay complex

We now construct a Delaunay complex. Although we can do this for any standard abstract Voronoi complex, we mainly focus on the two situations in which we know that a standard abstract Voronoi complex exists.

Definition 6.1. Let \mathcal{P} be a symmetric cocompact arrangement of k -dimensional flats with a standard abstract Voronoi complex \mathcal{V} . For any vertex $v \in \mathcal{V}$ labeled $\{\Pi_1, \dots, \Pi_{n+1}\}$, the link is the boundary of an n -dimensional simplex whose vertices are labeled $\{\Pi_1\}, \dots, \{\Pi_{n+1}\}$ and whose faces have labels equal to the union of their vertices' labels. Within this simplex, we draw a spine and label each cell of the spine with the same label as the simplex face to which it is dual. There is a label-preserving and orientation-preserving homeomorphism between this simplex and a closed regular neighborhood of v .

We construct an abstract complex \mathcal{D} by taking for each vertex $v \in \mathcal{V}$ an n -dimensional simplex (labeled as above), truncating the vertices from the simplex and performing certain face identifications between different truncated simplices (described below). A point that lies along a truncation face is regarded as having the same label as it did before the truncation. Let F be a j -dimensional face in \mathcal{V} and let the vertices of its boundary be v_1, \dots, v_m . Linearly identify the m different truncated simplices corresponding to each of the vertices v_1, \dots, v_m along their $(n - j)$ -dimensional faces labeled with the same label as F . For each face F in \mathcal{V} ,

perform these identifications. This produces an abstract complex \mathcal{D} that we call the drilled Delaunay complex.

Proposition 6.2. *The drilled Delaunay complex \mathcal{D} produces an n -dimensional manifold with boundary. For $k < n - 1$, each boundary component is homeomorphic to $\mathbb{R}^k \times S^{n-k-1}$, while for $k = n - 1$, each boundary component is homeomorphic to \mathbb{R}^{n-1} . The drilled Delaunay complex is homeomorphic (in an orientation-preserving way) to a closed regular neighborhood of $S \subset \mathbb{H}^n$, the union of all faces of \mathcal{V} . Thus, the interior of the drilled Delaunay complex is homeomorphic to $\mathbb{H}^n - \bigcup_{\Pi \in \mathcal{P}} \Pi$.*

Proof. Each truncated simplex in \mathcal{D} is homeomorphic (in an orientation-preserving way) to a closed regular neighborhood of the corresponding vertex of \mathcal{V} . The interior of each truncated simplex is obviously locally homeomorphic to \mathbb{R}^n . For any j -dimensional nontruncation face of one or more truncated simplices, the gluing pattern forces the link to be combinatorially equivalent to the boundary of one of the $(n - j)$ -dimensional faces of \mathcal{V} , and thus the link is homeomorphic to S^{n-j-1} . Thus, a regular neighborhood of the interior of a j -dimensional nontruncation face of \mathcal{D} is homeomorphic to $B^{n-j} \times \mathbb{R}^j$, and so is locally homeomorphic to \mathbb{R}^n . Any point in the interior of a truncation face obviously has a neighborhood homeomorphic to an n -dimensional half-ball. Any point that lies on the boundary of a truncation face also lies in the boundary of some nontruncation face. We've already determined that the link of a nontruncation face is a sphere, so such a point has a neighborhood homeomorphic to an n -dimensional half-ball.

By construction, the union of the spines of the truncated simplices is homeomorphic to S . Because \mathcal{V} is standard, we could construct a closed regular neighborhood of S that could be viewed as a (topological) complex of truncated simplices with a gluing pattern identical to that of \mathcal{D} . Thus, \mathcal{D} is homeomorphic to a closed regular neighborhood of $S \subset \mathbb{H}^n$. Finally, each boundary component of a closed regular neighborhood of S is homeomorphic to a component of the boundary of some n -dimensional cell of \mathcal{V} . The boundary of any n -dimensional cell of \mathcal{V} is homeomorphic to $S^{n-k-1} \times \mathbb{R}^k$. □

Definition 6.3. Define a map $\psi : \mathcal{D} \rightarrow \mathbb{H}^n$ by declaring that restricted to any one truncated simplex of \mathcal{D} , ψ is the map given by Theorem 4.12, using the flats of the simplex's label. Where different simplices intersect, the respective maps agree, so ψ is continuous.

We define an equivalence relation on \mathcal{D} to collapse the boundary.

Definition 6.4. For points x and y in the boundary of \mathcal{D} , we say that $x \sim y$ if there is a path γ from x to y lying entirely within the boundary of \mathcal{D} and such that ψ is constant along γ . Points not on the boundary of \mathcal{D} are equivalent only to themselves.

Proposition 6.5. *Every \sim equivalence class along the boundary of \mathfrak{D} is compact.*

Proof. Since ψ is continuous, the preimage of a point is closed. The equivalence classes are components of the preimages through ψ of points in $\bigcup_{\Pi \in \mathcal{P}} \Pi$. The components of a closed subspace are themselves closed.

Since there are only finitely many Γ -orbits of vertices in the standard abstract Voronoi complex, there are also only finitely many Γ -orbits of truncated simplices in \mathfrak{D} . The image under ψ of any one truncated simplex is a Delaunay cell. Since Delaunay cells are bounded, any point $x \in \bigcup_{\Pi \in \mathcal{P}} \Pi$ can lie in only finitely many of them. Thus, the preimage in \mathfrak{D} of any point must be bounded. \square

Proposition 6.6. *For $n = 3, k = 1$, if C is a \sim equivalence class along the boundary of \mathfrak{D} that contains no vertices of truncation faces, then C is homeomorphic to S^1 .*

Proof. By Theorem 4.12, the portion of C that lies within any one truncation face can be chosen to be either a vertex or an affine subspace of dimension at least 1. Since the truncation faces are 2-dimensional, the portion of C lying within any one truncation face is either a vertex, a line segment, or the entire truncation face. By assumption, C contains no truncation face vertices, so C must be a union of line segments. Any one line segment has two endpoints, lying in the interiors of truncation face edges. Thus, each endpoint also lies in some other truncation face, where it must be the endpoint of a unique different line segment. Thus, C is a compact connected 1-dimensional manifold, and so it is homeomorphic to S^1 . \square

Proposition 6.7. *For $n = 3$ and $k = 1$, if C is a \sim equivalence class along the boundary of \mathfrak{D} that contains a vertex of a truncation face, then a regular neighborhood of C within the boundary of \mathfrak{D} is homeomorphic to a punctured sphere.*

Proof. Each component of the boundary of \mathfrak{D} is homeomorphic to a doubly infinite cylinder $S^1 \times \mathbb{R}$. The equivalence class C is a connected compact subset of $S^1 \times \mathbb{R}$, consisting of a union of a finite number of vertices, edges, and triangles. Let U be a closed regular neighborhood of C . Then each component of the boundary of U is a closed 1-dimensional manifold, so it is a circle. There are either one or two unbounded components of $(S^1 \times \mathbb{R}) - U$, so there are either one or two circles in the boundary of U that meet unbounded components of $(S^1 \times \mathbb{R}) - U$. If there is only one such circle, then it bounds a disk within $S^1 \times \mathbb{R}$, so U is contained within a disk. If there are two such circles, then they bound an annulus within $S^1 \times \mathbb{R}$, so U is contained within an annulus. Either way, U is contained within a punctured sphere and has circles as boundary components, so U is also homeomorphic to a punctured sphere. \square

Proposition 6.8. *For $n = 3$ and $k = 1$, \mathfrak{D}/\sim is a 3-dimensional manifold.*

Proof. Within the interior of \mathfrak{D} , \sim identifies points only with themselves. Thus, the quotient of the interior of \mathfrak{D} by \sim is still the interior of \mathfrak{D} , which is homeomorphic

to $\mathbb{H}^3 - \bigcup_{\Pi \in \mathcal{P}} \Pi$. All remaining points of \mathcal{D}/\sim are of the types described in the previous two propositions. A regular neighborhood of such a point is a regular neighborhood of the corresponding equivalence class modulo \sim .

For an equivalence class of the type described in Proposition 6.6, a small enough neighborhood can be chosen so that the only boundary equivalence classes it contains are also of the same type. Thus, a regular neighborhood in \mathcal{D} is homeomorphic to $S^1 \times \{(x, y) \in B^2 : y \geq 0\}$ with \sim collapsing $S^1 \times \{(x, 0) \in B^2\}$ along the S^1 coordinate. Thus, a regular neighborhood of such a point in \mathcal{D}/\sim is homeomorphic to B^3 .

For an equivalence class C of the type described in Proposition 6.7, a small enough closed neighborhood U in the boundary of \mathcal{D} can be chosen so that all other boundary equivalence classes that it contains are of the type described in Proposition 6.6, and so $U - C$ is homeomorphic to a Cartesian product of S^1 with some disjoint union of half-open line segments, in which \sim collapses along the S^1 coordinate. The boundary of a regular neighborhood of C in \mathcal{D} is homeomorphic to two copies of U , U_1 and U_2 glued along their boundaries, which are disjoint unions of circles. We take U_1 to be the copy of U that lies in the boundary of \mathcal{D} and U_2 to be the copy that lies in the interior of \mathcal{D} (except where it's glued to U_1).

After the quotient by \sim , U_1/\sim contains only interior points (except along its intersection with U_2), while U_2/\sim is the boundary of the regular neighborhood. Each component of the intersection of U_2 with U_1 is a circle, which \sim identifies to a point. By Proposition 6.7, U_2 is homeomorphic to a sphere with a finite number of open disks removed, so collapsing each boundary component of U_2 to a point produces a sphere.

Thus, the boundary of sufficiently small regular neighborhoods of $C/\sim \in \mathcal{D}/\sim$ are spheres, so at C/\sim , \mathcal{D}/\sim is locally homeomorphic to \mathbb{R}^3 . □

Proposition 6.9. *For $n = 3$ and $k = 1$, \mathcal{D}/\sim is homeomorphic to \mathbb{H}^3 in a Γ -equivariant fashion, so $(\mathcal{D}/\sim)/\Gamma$ is homeomorphic to \mathbb{H}^3/Γ .*

Proof. Each boundary component of \mathcal{D} is homeomorphic to a doubly infinite cylinder, so any sufficiently small regular neighborhood of a boundary component has a boundary consisting of two doubly infinite cylindrical components. The quotient of such a neighborhood by \sim is a manifold whose boundary is a doubly infinite cylindrical component, as the cylinder lying along the boundary of \mathcal{D} is collapsed to interior points. Thus, for each component of the boundary of \mathcal{D} , every sufficiently small regular neighborhood of its quotient by \sim has cylindrical boundary. Since \mathcal{D}/\sim is a manifold, it must be the case that every such regular neighborhood is homeomorphic to a solid doubly infinite cylinder. As the interior of \mathcal{D} is homeomorphic to $\mathbb{H}^3 - \bigcup_{\Pi \in \mathcal{P}} \Pi$ and each boundary component of \mathcal{D} is a doubly infinite cylinder, we have that \mathcal{D}/\sim is homeomorphic to \mathcal{D} with a doubly infinite solid

cylinder glued to each boundary component. Up to homeomorphism, there's only one way to do this, which is to attach a regular neighborhood of each $\Pi \in \mathcal{P}$ to $\mathbb{H}^3 - \bigcup_{\Pi \in \mathcal{P}} \Pi$.

The action of Γ on \mathcal{D} and \sim is compatible with the action of Γ on \mathbb{H}^3 , so \mathcal{D}/\sim is homeomorphic to \mathbb{H}^3 in a Γ -equivariant fashion. \square

Theorem 6.10. *For $n = 3, k = 1$, the map $\psi : \mathcal{D} \rightarrow \mathbb{H}^3$ induces a degree-1 map from $(\mathcal{D}/\sim)/\Gamma$ to \mathbb{H}^3/Γ .*

Proof. On any given \sim equivalence class, ψ is constant. Thus, ψ induces a map $\psi' : \mathcal{D}/\sim \rightarrow \mathbb{H}^3$. There is a Γ -action on \mathcal{D}/\sim and ψ is Γ -equivariant, so ψ' is Γ -equivariant. The homeomorphism $\chi : \mathcal{D}/\sim \rightarrow \mathbb{H}^3$ produced by Proposition 6.9 is also Γ -equivariant, so the straight line homotopy between ψ' and χ is also Γ -equivariant. Thus, the induced maps $\psi'' : (\mathcal{D}/\sim)/\Gamma \rightarrow \mathbb{H}^3/\Gamma$ and $\chi' : (\mathcal{D}/\sim)/\Gamma \rightarrow \mathbb{H}^3/\Gamma$ are homotopic, and so have the same degree.

Restricted to the interior of \mathcal{D} (where \sim equivalence classes are single points), χ is orientation-preserving. Thus, the degree of χ' and ψ' must be 1. \square

The map ψ is constructed using the map from Theorem 4.12, which is sometimes orientation-reversing. Also, the Voronoi decomposition for lines in \mathbb{H}^3 can have multiple vertices with the same label. Thus, there might be multiple truncated simplices in \mathcal{D} that are mapped to the same cell in \mathbb{H}^3 .

Corollary 6.11. *Counting cells by their multiplicities and orientations, for $n = 3, k = 1$, almost every point of \mathbb{H}^3 is in a total of one open 3-dimensional Delaunay cell.*

Proof. At a regular value $x \in \mathbb{H}^3/\Gamma$ of ψ'' , we can compute the degree of ψ'' by summing the signed orientation of ψ'' over $\psi''^{-1}(x)$. The only points of \mathbb{H}^3/Γ that possibly aren't regular values are those that lie either in one of the flats of the arrangement \mathcal{P} or within some closed Delaunay cell of dimension less than n . \square

We now prove a similar result for the case $k = n - 1$. This is a generalization of Marshall and Martin's 2-dimensional result [2003].

Proposition 6.12. *For $k = n - 1$, the map ψ produced by Definition 6.3 maps the interior of \mathcal{D} to $\mathbb{H}^n - \bigcup_{\Pi \in \mathcal{P}} \Pi$.*

Proof. Any point in the interior of \mathcal{D} is either interior to some T in \mathcal{D} or on a nontruncation face of some T in \mathcal{D} , and so is mapped by ψ to a point in some closed ideal Delaunay cell. Let the label for T be $\{\Pi_1, \dots, \Pi_{n+1}\}$. Let v be the Voronoi vertex corresponding to T . The flats Π_1, \dots, Π_{n+1} separate \mathbb{H}^n into various connected components. Since v is equidistant from the flats Π_1, \dots, Π_{n+1} and is not closer to any other flat in \mathcal{P} , from v one can draw a line segment to any flat Π_i ($1 \leq i \leq n + 1$) without crossing any other flats in \mathcal{P} . Thus, v must lie in the

unique component U of $\mathbb{H}^n - \bigcup_{\Pi \in \mathcal{P}} \Pi_i$ that contains all of Π_1, \dots, Π_{n+1} within its boundary.

Let \mathbf{x} be a point that is not in the closure of U . Then there is some $\Pi \in \mathcal{P}$ that separates \mathbf{x} from U . The flat Π then separates \mathbf{x} from either all of $\pi_1(\mathbf{x}), \dots, \pi_{n+1}(\mathbf{x})$ or all but one of $\pi_1(\mathbf{x}), \dots, \pi_{n+1}(\mathbf{x})$ (if Π is one of Π_1, \dots, Π_{n+1}), so the relative interior of the convex hull of $\pi_1(\mathbf{x}), \dots, \pi_{n+1}(\mathbf{x})$ is on the opposite side of Π from \mathbf{x} . Then \mathbf{x} does not lie in the closed ideal Delaunay cell. Thus, the closed Delaunay cell lies within the closure of U .

Let \mathbf{x} be a point in the closed ideal Delaunay cell and suppose that there is some $\Pi \in \mathcal{P}$ that contains \mathbf{x} . Then \mathbf{x} lies within the relative interior of the convex hull of some non singleton subset of $\pi_1(\mathbf{x}), \dots, \pi_{n+1}(\mathbf{x})$; without loss of generality, say the subset is $\pi_1(\mathbf{x}), \dots, \pi_m(\mathbf{x})$ and that Π is not one of Π_1, \dots, Π_m . However, since Π is $(n - 1)$ -dimensional, contains none of $\pi_1(\mathbf{x}), \dots, \pi_m(\mathbf{x})$, and yet intersects the interior of their convex hull, it separates some of them from each other. However, that is impossible because we know that there is a connected component U of $\mathbb{H}^n - \bigcup_{\Pi \in \mathcal{P}} \Pi$ that contains all of Π_1, \dots, Π_{n+1} within its boundary. \square

Theorem 6.13. *Counting cells by their multiplicities and orientations, for $k = n - 1$, almost every point of \mathbb{H}^n is in a total of one open n -dimensional Delaunay cell.*

Proof. We cut \mathbb{H}^n / Γ along the compact connected $(n - 1)$ -dimensional surface $(\bigcup_{\Pi \in \mathcal{P}} \Pi) / \Gamma$ to produce a manifold N with two boundary components, each isometric to $(\bigcup_{\Pi \in \mathcal{P}} \Pi) / \Gamma$.

The interior of \mathcal{D} is homeomorphic in a Γ -equivariant orientation-preserving fashion to $\mathbb{H}^n - \bigcup_{\Pi \in \mathcal{P}} \Pi$. Restricted to the interior of \mathcal{D} , ψ is also a Γ -equivariant map to $\mathbb{H}^n - \bigcup_{\Pi \in \mathcal{P}} \Pi$. Thus, there is a Γ -equivariant homotopy between the two maps. For any one connected component of $\mathbb{H}^n - \bigcup_{\Pi \in \mathcal{P}} \Pi$, each boundary component of its closure is some $\Pi \in \mathcal{P}$. The quotient of this flat by its stabilizer in Γ is isometric to $(\bigcup_{\Pi \in \mathcal{P}} \Pi) / \Gamma$.

Each boundary component of \mathcal{D} is homeomorphic to an $(n - 1)$ -dimensional flat. The quotient of such a flat by its stabilizer in Γ will be homeomorphic to $(\bigcup_{\Pi \in \mathcal{P}} \Pi) / \Gamma$.

Then $\psi : \mathcal{D} \rightarrow \mathbb{H}^n$ induces a map (of manifolds with boundary) $\psi' : \mathcal{D} / \Gamma \rightarrow N$. This map ψ' is homotopic to a homeomorphism, and so has degree 1, as a map of compact manifolds with boundary.

As in the proof of Corollary 6.11, each regular value of ψ' must be in a total of one Delaunay cell, when counting cells according to multiplicities and orientations. The only points in N that possibly aren't regular are either in some $\Pi \in \mathcal{P}$ or in some open ideal Delaunay cell of dimension less than n , so almost every point in N is a regular value.

Since \mathbb{H}^n/Γ is just N with its two boundary components glued together, we have that almost every point of \mathbb{H}^n/Γ is in one Delaunay cell. Lifting the map, almost every point of \mathbb{H}^n is in one Delaunay cell. \square

References

- [Böröczky 1978] K. Böröczky, “Packing of spheres in spaces of constant curvature”, *Acta Math. Acad. Sci. Hungar.* **32**:3-4 (1978), 243–261. MR 80h:52014 Zbl 0422.52011
- [Dowty 2000] J. G. Dowty, *Ortholengths and hyperbolic Dehn surgery*, thesis, University of Melbourne, 2000, Available at <http://repository.unimelb.edu.au/10187/3534>.
- [Gabai et al. 2009] D. Gabai, R. Meyerhoff, and P. Milley, “Minimum volume cusped hyperbolic three-manifolds”, *J. Amer. Math. Soc.* **22**:4 (2009), 1157–1215. MR 2011a:57031 Zbl 1204.57013
- [Gabai et al. 2010] D. Gabai, R. Meyerhoff, and P. Milley, “Mom technology and hyperbolic 3-manifolds”, pp. 84–107 in *In the tradition of Ahlfors–Bers, V* (Newark, NJ, 2008), edited by M. Bonk et al., Contemp. Math. **510**, Amer. Math. Soc., Providence, RI, 2010. MR 2011f:57026
- [Marshall and Martin 2003] T. H. Marshall and G. J. Martin, “Packing strips in the hyperbolic plane”, *Proc. Edinb. Math. Soc. (2)* **46**:1 (2003), 67–73. MR 2004a:52032 Zbl 1032.52007
- [Milley 2009] P. Milley, “Minimum volume hyperbolic 3-manifolds”, *J. Topol.* **2**:1 (2009), 181–192. MR 2010d:57018 Zbl 1165.57016
- [Przeworski 2006] A. Przeworski, “A universal upper bound on density of tube packings in hyperbolic space”, *J. Differential Geom.* **72**:1 (2006), 113–127. MR 2007b:57030 Zbl 1098.52005
- [Przeworski 2012] A. Przeworski, “An upper bound on density for packings of collars about hyperplanes in \mathbb{H}^n ”, *Geom. Dedicata Online First* (2012), 1–21.
- [Ratcliffe 1994] J. G. Ratcliffe, *Foundations of hyperbolic manifolds*, Graduate Texts in Mathematics **149**, Springer, New York, 1994. MR 95j:57011 Zbl 0809.51001
- [Rogers 1958] C. A. Rogers, “The packing of equal spheres”, *Proc. London Math. Soc. (3)* **8** (1958), 609–620. MR 21 #847 Zbl 0085.03302

Received August 10, 2011.

ANDREW PRZEWORSKI
 DEPARTMENT OF MATHEMATICS
 COLLEGE OF CHARLESTON
 66 GEORGE ST.
 CHARLESTON, SC 29424
 przeworskia@cofc.edu
<http://przeworskia.people.cofc.edu>

PACIFIC JOURNAL OF MATHEMATICS

<http://pacificmath.org>

Founded in 1951 by
E. F. Beckenbach (1906–1982) and F. Wolf (1904–1989)

EDITORS

V. S. Varadarajan (Managing Editor)
Department of Mathematics
University of California
Los Angeles, CA 90095-1555
pacific@math.ucla.edu

Vyjayanthi Chari
Department of Mathematics
University of California
Riverside, CA 92521-0135
chari@math.ucr.edu

Darren Long
Department of Mathematics
University of California
Santa Barbara, CA 93106-3080
long@math.ucsb.edu

Sorin Popa
Department of Mathematics
University of California
Los Angeles, CA 90095-1555
popa@math.ucla.edu

Robert Finn
Department of Mathematics
Stanford University
Stanford, CA 94305-2125
finn@math.stanford.edu

Jiang-Hua Lu
Department of Mathematics
The University of Hong Kong
Pokfulam Rd., Hong Kong
jhlu@maths.hku.hk

Jie Qing
Department of Mathematics
University of California
Santa Cruz, CA 95064
qing@cats.ucsc.edu

Kefeng Liu
Department of Mathematics
University of California
Los Angeles, CA 90095-1555
liu@math.ucla.edu

Alexander Merkurjev
Department of Mathematics
University of California
Los Angeles, CA 90095-1555
merkurev@math.ucla.edu

Jonathan Rogawski
Department of Mathematics
University of California
Los Angeles, CA 90095-1555
jonr@math.ucla.edu

PRODUCTION

pacific@math.berkeley.edu

Silvio Levy, Scientific Editor

Matthew Cargo, Senior Production Editor

SUPPORTING INSTITUTIONS

ACADEMIA SINICA, TAIPEI
CALIFORNIA INST. OF TECHNOLOGY
INST. DE MATEMÁTICA PURA E APLICADA
KEIO UNIVERSITY
MATH. SCIENCES RESEARCH INSTITUTE
NEW MEXICO STATE UNIV.
OREGON STATE UNIV.

STANFORD UNIVERSITY
UNIV. OF BRITISH COLUMBIA
UNIV. OF CALIFORNIA, BERKELEY
UNIV. OF CALIFORNIA, DAVIS
UNIV. OF CALIFORNIA, LOS ANGELES
UNIV. OF CALIFORNIA, RIVERSIDE
UNIV. OF CALIFORNIA, SAN DIEGO
UNIV. OF CALIF., SANTA BARBARA

UNIV. OF CALIF., SANTA CRUZ
UNIV. OF MONTANA
UNIV. OF OREGON
UNIV. OF SOUTHERN CALIFORNIA
UNIV. OF UTAH
UNIV. OF WASHINGTON
WASHINGTON STATE UNIVERSITY

These supporting institutions contribute to the cost of publication of this Journal, but they are not owners or publishers and have no responsibility for its contents or policies.

See inside back cover or pacificmath.org for submission instructions.

The subscription price for 2012 is US \$420/year for the electronic version, and \$485/year for print and electronic. Subscriptions, requests for back issues from the last three years and changes of subscribers address should be sent to Pacific Journal of Mathematics, P.O. Box 4163, Berkeley, CA 94704-0163, U.S.A. Prior back issues are obtainable from Periodicals Service Company, 11 Main Street, Germantown, NY 12526-5635. The Pacific Journal of Mathematics is indexed by Mathematical Reviews, Zentralblatt MATH, PASCAL CNRS Index, Referativnyi Zhurnal, Current Mathematical Publications and the Science Citation Index.

The Pacific Journal of Mathematics (ISSN 0030-8730) at the University of California, c/o Department of Mathematics, 969 Evans Hall, Berkeley, CA 94720-3840, is published monthly except July and August. Periodical rate postage paid at Berkeley, CA 94704, and additional mailing offices. POSTMASTER: send address changes to Pacific Journal of Mathematics, P.O. Box 4163, Berkeley, CA 94704-0163.

PJM peer review and production are managed by EditFLOW™ from Mathematical Sciences Publishers.

PUBLISHED BY PACIFIC JOURNAL OF MATHEMATICS

at the University of California, Berkeley 94720-3840

A NON-PROFIT CORPORATION

Typeset in L^AT_EX

Copyright ©2012 by Pacific Journal of Mathematics

PACIFIC JOURNAL OF MATHEMATICS

Volume 258 No. 1 July 2012

On the complexity of sails	1
LUKAS BRANTNER	
Construction of Lagrangian submanifolds in $\mathbb{C}\mathbb{P}^n$	31
QING CHEN, SEN HU and XIAOWEI XU	
Semisimple tunnels	51
SANGBUM CHO and DARRYL MCCULLOUGH	
Degenerate two-boundary centralizer algebras	91
ZAJJ DAUGHERTY	
Heegaard genera in congruence towers of hyperbolic 3-manifolds	143
BOGWANG JEON	
The Heisenberg ultrahyperbolic equation: The basic solutions as distributions	165
ANTHONY C. KABLE	
Rational Seifert surfaces in Seifert fibered spaces	199
JOAN E. LICATA and JOSHUA M. SABLOFF	
Delaunay cells for arrangements of flats in hyperbolic space	223
ANDREW PRZEWORSKI	



0030-8730(201209)258:1;1-2