Algorithms for the Epstein-Penner decomposition of a cusped surface

Sampson Wong

An essay submitted in partial fulfillment of
the requirements for the degree of
B.Sc. (Honours)

Pure Mathematics
University of Sydney

October 2014

CONTENTS

Introduction

The Epstein-Penner decomposition is an elegant yet powerful construction in the study of non-compact finite volume hyperbolic manifolds. The Epstein-Penner construction first arose in the study of Teichmüller space and its cell decomposition, building on earlier work by Harer [**Har86**] and by Bowditch and Epstein [**BE88**]. In [**Pen87**], Penner uses the Epstein-Penner construction to assign an ideal cell decomposition to each point of Teichmüller space, which remarkably induces a cell decomposition of Teichmüller space, the proof of which is extremely difficult.

The Epstein-Penner construction has since found many uses in the study of compact orientable surfaces [**Ush99**], hyperbolic 3-manifolds [**Yos01**] [**HRS12**] and 4-manifolds [**KM13**] to name a few examples.

An efficient algorithm to compute the Epstein-Penner decomposition of a non-compact hyperbolic manifold was first given by Weeks in [**Wee93**]. Weeks' algorithm plays an extremely important role in the systematic study of hyperbolic 3-manifolds [**CHW99**], and is included in the popular software packages *SnapPea* and *Regina*. Although Weeks' algorithm is not a true algorithm and only a heuristic procedure for dimensions 3 and greater, it is very reliable in practice.

In the recent paper by Cooper-Long [**CL13**], the Epstein-Penner construction was generalised to non-compact, finite volume, strictly convex projective manifolds, building on earlier work by Cooper, Long and Tillman [**CLT11**]. This raises the following open problems: Is there an efficient algorithm to compute the Epstein-Penner decomposition in this case? Do these cell decompositions induce a cell decomposition of its Teichmüller space in the case of non-compact strictly convex projective surfaces? Which results about the Epstein-Penner decomposition generalise to the strictly convex projective setting?

This paper attempts to answer the first question by proposing a novel algorithm for the Epstein-Penner decomposition. Unlike its predecessors, it is not restricted to hyperbolic geometry and is able to compute the Epstein-Penner decomposition of strictly convex projective surfaces. Also, since hyperbolic geometry is a subgeometry of projective geometry, the proposed algorithm is still applicable to cusped hyperbolic manifolds. The algorithm

may be a useful tool for further study of strictly convex projective manifolds, for example, to study the second problem.

In Chapter 1, some basic theory of hyperbolic geometry is reviewed and in Chapter 2 the construction of the Epstein-Penner decomposition is explained. In Chapters 3 and 4, analogous results are established in the case of strictly convex projective manifolds, including Cooper-Long's construction. Finally, in Chapter 5 the new algorithm is presented with proof of correctness for strictly convex projective surfaces.

CHAPTER 1

Background - Hyperbolic Geometry

1.1. Hyperboloid Model of Hyperbolic Space

Define *Minkowski space* to be the real vector space $\mathbb{R}^3$ with the quadratic form $\langle \cdot, \cdot \rangle$ defined by

$$\langle x, y \rangle = -x_1 y_1 + x_2 y_2 + x_3 y_3.$$

The form $\langle \cdot, \cdot \rangle$ is called the *Minkowski Inner Product*.

Vectors with inner product $\langle x, x \rangle = 0$ are called *light-like* vectors. A light-like vector is called positive (resp. negative) if $x_1 > 0$ (resp. $x_1 < 0$). The set of all positive light-like vectors forms the *positive light-cone*, which we will denote by L^+.

The vectors satisfying $\langle x, x \rangle = -1$ form a two sheeted hyperboloid. The hyperboloid model of the hyperbolic plane is given by the upper sheet

$$H = \{x \in M : \langle x, x \rangle = -1, x_1 \geq 1\},$$

imbued with the Riemannian metric inherited from the quadratic form $\langle \cdot, \cdot \rangle$. Points at infinity are represented by rays on the light-cone L^+. The isometries of H are the 3×3 matrices preserving the quadratic form $\langle \cdot, \cdot \rangle$ and that do not exchange the two sheets of the hyperboloid. This group of isometries, which we will call $SO^+(1, 2)$, are the matrices $A \in GL(3, \mathbb{R})$ such that $\det A = 1$, the top left entry is positive and $A^t J A = J$ where $J = \text{diag}(-1, 1, 1)$ [**Rat06**].

1.2. Upper Half-plane Model

The *upper half-plane model* of hyperbolic geometry is the subset $\mathbb{H}^2 = \{x + iy \in \mathbb{C} : y > 0\}$ of the complex plane with the metric

$$ds = \frac{\sqrt{dx^2 + dy^2}}{y}.$$

The geodesics of the model are either circular arcs perpendicular to the real axis or straight vertical lines ending on the real axis.

The *special linear group* $SL(2, \mathbb{R})$ is the group of 2×2 matrices with determinant 1. Quotienting the special linear group by its centre gives the

projective special linear group

$$PSL(2,\mathbb{R}) = SL(2,\mathbb{R})/\{\pm \left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)\}.$$

Elements in $PSL(2,\mathbb{R})$ act on the upper half of the complex plane $\mathbb{H}^2$ via linear fractional transformations.

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot z = \frac{az+b}{cz+d}$$

Note that the matrix is only well defined up to sign, however, the sign of the matrix does not affect the transformation. Hence, the action is well defined. These maps are the orientation preserving isometries of the upper half-plane $\mathbb{H}$ and are called *Möbius Transformations*. Futhermore, the composition of Möbius transformations is equivalent to matrix multiplication in $PSL(2,\mathbb{R})$.

1.3. Poincaré Disk Model

The *Poincaré Disk*, denoted by $\mathbb{D}$, is defined to be the disk $\{x \in \mathbb{R}^2 : |x| < 1\}$ with the metric

$$ds = \frac{\sqrt{dx^2+dy^2}}{1-x^2-y^2}.$$

There are two types of geodesics of the Poincaré Disk Model. The first type is circular arcs whose endpoints are perpendicular to the boundary $\partial\mathbb{D} = \{x \in \mathbb{R}^2 \, |x| = 1\}$, and the second type is the diameters of the disk. The second case may be considered a special case of the first if we allow circular arcs to have infinite radius.

There is a bijection between the upper half-plane $\mathbb{H}^2$ and the Poincaré Disk $\mathbb{D}$ given by the map $h : \mathbb{H}^2 \to \mathbb{D}$,

$$h(z) = \frac{z-i}{iz-1}.$$

Moreover, h is an isometry. If $A \in PSL(2,\mathbb{R})$ is an orientation preserving isometry $\mathbb{H}^2 \to \mathbb{H}^2$, then hAh^{-1} is also an orientation preserving isometry of $\mathbb{D}$. Hence, $A \in PSL(2,\mathbb{R})$ acts on the Poincaré disk via the orientation preserving isometries hAh^{-1}.

Both the Poincaré model and the upper half-plane model are conformal representations of the hyperbolic plane, meaning that the angles between lines in the hyperbolic space are same as the Euclidean ones in the model.

The Poincaré model and the hyperboloid model are related projectively. Take the hyperboloid H defined in section 1.1 and let the Poincaré Disk be the set $\{x = (0, x_2, x_3) : |x| < 1\} \subset \mathbb{R}^3$. Radial projection from the projection point $(-1, 0, 0)$ relates the two models, and is in fact an isometry. Moreover, this projection lifts the action of $PSL(2,\mathbb{R})$ on $\mathbb{D}$ to an action on

H since the projection is an isometry. So $PSL(2,\mathbb{R})$ acts as orientation preserving isometries on H, in fact, all orientation preserving isometries of H arise this way. The explicit action of $PSL(2,\mathbb{R})$ on H will be discussed in section 1.6.

1.4. Klein Model

The *Projective* or *Klein model* of the hyperbolic plane is given by the set

$$D = \{x \in \mathbb{R}^2 : |x| < 1\},$$

together with a hyperbolic metric given by the following formula: Let p and q be distinct points in D and let the line pq intersect the boundary of the Klein model ∂D at points a and b so that the points are, in order, a, p, q, b. Then,

$$d(p,q) = \frac{1}{2}\log\frac{|qa||bp|}{|pa||qb|},$$

where $|\cdot|$ denotes the Euclidean distance. Geodesics in the Klein model are the Euclidean ones.

An isometry between the Poincaré Disk and the Klein model is given by $k : \mathbb{D} \to D, u \mapsto \frac{2u}{1+|u|^2}$. Under this transformation, points on the boundary are fixed, and a curved geodesic with endpoints $a, b \in \partial\mathbb{D}$ is mapped to a straight geodesic with the same endpoints $a, b \in D$.

The isometry between D and H is given by radial projection, this time D is the set $\{x = (1, x_2, x_3) : x_2^2 + x_3^2 < 1\} \in \mathbb{R}^3$ and the centre of projection is the origin. This projection of the Klein model to the hyperboloid model is an important step in the Epstein-Penner Construction, which will be discussed in Chapter 2.

The orientation preserving isometries of the Klein model D are the orientation preserving isometres of $\mathbb{D}$ conjugated by the isometry k, hence we can identify the orientation preserving isometries of D with the group elements in $PSL(2,\mathbb{R})$. The isometries of the Klein model are also the orientation preserving isometries of H conjugated by the projection ρ, and can be identified with the elements of $SO^+(1,2)$.

1.5. Horocycles

In the upper half-plane model, horocycles are either a circle tangent to real axis or a horizontal line $l_a = \{x + ia : x \in \mathbb{R}\}$, $a > 0$. A horocycle tangent to the real axis at $x \in \mathbb{R}$ is said to be *centred* at x, whereas the horizontal horoballs are all centred at ∞. The metric on the horocycle l_a is the Euclidean distance with a scaling factor of $\frac{1}{a}$.

In the Poincaré disk model, horocycles are circles tangent to the boundary.

In the hyperboloid model, a horocycle is the intersection of H with a plane in $\mathbb{R}^3$ whose normal vector is light-like. The horocycle h_w may be naturally associated with this light like vector $w \in L^+$ by the Minkowski inner product: $h_w = \{x \in H : \langle w, x\rangle = -1\}$. Hence, the closer the horocycle h_w is to infinity, the closer its light-like representative w is to the origin.

Lemma 1.1. *The horocycle l_a in the upper half-plane model corresponds to a horocycle in the Poincaré Disk model with radius $\frac{1}{1+a}$.*

Proof. The isometry between the upper half-plane model and the Poincaré disk is given by

$$h(z) = \frac{z-i}{iz-1},$$

so the centre of horocycle l_a is sent to $h(\infty) = -i$. Then $h(l_a)$ is a circle which lies inside the disk $|z| \leq 1$ and tangent to it at $-i$. The point diametrically opposite to $h(\infty)$ in the Poincare model is $h(ia) = \frac{i(1-a)}{1+a}$. The diameter lies on the imaginary axis and has length $\frac{2}{1+a}$, hence the horocycle l_a corresponds to a horocycle in $\mathbb{D}$ with Euclidean radius $\frac{1}{1+a}$. □

Lemma 1.2. *Let h_w be a horocycle in the hyperboloid model with associated light like vector $w \in L^+$ having first coordinate, or height, a. Then h_w corresponds to a horocycle in the Poincaré Disk model with Euclidean radius $\frac{1}{1+a}$.*

Proof. Without loss of generality let $w = (a, a, 0)$. Since ray w points towards the centre of horoball h_w, the projection of h_w onto the Poincare model will be a horoball centred at $(0, 1, 0)$. The point diametrically opposite will be of the form $(0, k, 0)$, and projecting this point onto the hyperboloid model gives a point of the form $(x, y, 0)$ (see Figure 1.5.1).

The point $(x, y, 0)$ must be on the hyperboloid and the horoball h_w, therefore $-x^2 + y^2 = -1$ and $-ax + ay = -1$. Solving the equations simultaneously gives $x = \frac{1}{2}(a + \frac{1}{a})$ and $y = \frac{1}{2}(a - \frac{1}{a})$, and projecting the point back onto the Poincaré disk gives the point $(0, \frac{2a}{a+1} - 1, 0)$ diametrically opposite ot $(0, 1, 0)$ on the horoball. Hence, the Euclidean radius of the horoball in the Poincaré model is $\frac{1}{2}(2 - \frac{2a}{a+1}) = \frac{1}{1+a}$. □

1.6. Isometries of hyperbolic space

Let $A \in PSL(2, \mathbb{R})$ be a Möbius Transformation. If A is not the identity transformation, then by considering its Jordan Normal form, A is conjugate to $\pm \left(\begin{smallmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{smallmatrix}\right)$, $\pm \left(\begin{smallmatrix} a & 0 \\ 0 & 1/a \end{smallmatrix}\right)$, or $\pm \left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$, and is called *elliptic*, *hyperbolic* or *parabolic* respectively. Elliptic transformations have $|\operatorname{tr}(A)| < 2$ and fixes exactly one point in $\mathbb{H}^2$. Hyperbolic transformations have $|\operatorname{tr}(A)| > 2$ and

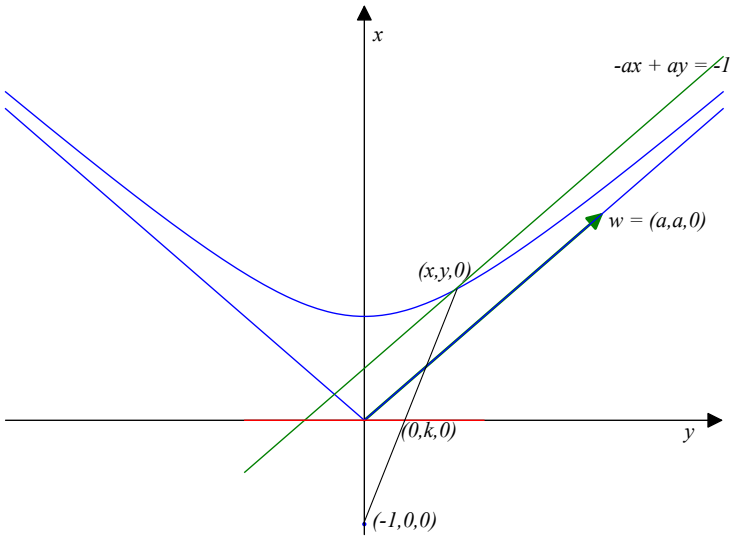

FIGURE 1.5.1. The projection of a horosphere with light-like representative $w = (a, a, 0)$ onto the Poincaré Disk.

fix two points on the boundary $\mathbb{R} \cup \{\infty\}$. Parabolic transformations have $|\operatorname{tr}(A)| = 2$ and fixes exactly one point on the boundary $\mathbb{R} \cup \{\infty\}$.

The group of orientation preserving isometries of hyperbolic space is $PSL(2, \mathbb{R})$, however, the orientation preserving isometries of H are also given by linear transformations preserving the bilinear form $\langle \cdot, \cdot \rangle$ in section 1.1. Indeed, the action of group elements of $PSL(2, \mathbb{R})$ on H is given in [**Pen87**] as follows.

We represent a point $x = (x_1, x_2, x_3) \in \mathbb{R}^3$ by the matrix $\left(\begin{smallmatrix} x_1+x_3 & x_2 \\ x_2 & x_1-x_3 \end{smallmatrix}\right)$. This identifies H with the set of symmetric matrices with determinant $x_1^2 - x_2^2 - x_3^2 = 1$ [todo]. The action of $A \in PSL(2, \mathbb{R})$ on this matrix is given by

$$\begin{pmatrix} x_1 + x_3 & x_2 \\ x_2 & x_1 - x_3 \end{pmatrix} \mapsto A^t \begin{pmatrix} x_1 + x_3 & x_2 \\ x_2 & x_1 - x_3 \end{pmatrix} A.$$

The action of $A = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ on $\mathbb{R}^3$ is calculated from the above expression.

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} \mapsto \begin{pmatrix} \frac{1}{2}(a^2 + b^2 + c^2 + d^2) & ac + bd & \frac{1}{2}(a^2 + b^2 - c^2 - d^2) \\ ab + cd & bc + ad & ab - cd \\ \frac{1}{2}(a^2 - b^2 + c^2 - d^2) & ac - bd & \frac{1}{2}(a^2 - b^2 - c^2 + d^2) \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix}$$

Let $\widetilde{A}$ be the 3×3 matrix shown above. Then $\widetilde{A} \in SO^+(1, 2)$ since $\det \widetilde{A} = 1$, the top left entry of $\widetilde{A}$ is positive and $\widetilde{A}^t J \widetilde{A} = J$, where $J = \operatorname{diag}(-1, 1, 1)$. Furthermore, the map $\alpha : PSL(2, \mathbb{R}) \to SO^+(1, 2), A \mapsto$

$\widetilde{A}$ is a group isomorphism. The map sends $\pm\left(\begin{smallmatrix}1&1\\0&1\end{smallmatrix}\right)$ to $\left(\begin{smallmatrix}1&1&0\\0&1&1\\0&0&1\end{smallmatrix}\right)$. and hence sends parabolic elements of $PSL(2,\mathbb{R})$ to parabolic elements of $SO^+(1,2)$ conjugate to the standard parabolic $\left(\begin{smallmatrix}1&1&0\\0&1&1\\0&0&1\end{smallmatrix}\right)$. In particular, the parabolics in $SO^+(1,2)$ form a single conjugacy class.

1.7. Hyperbolic surfaces

The Möbius Transformations $PSL(2,\mathbb{R})$ forms a topological group under composition and the *compact-open topology*. The compact-open topology of a group of isometries Γ on a topological space X is generated by the subsets of the form $B_{K,U} = \{A \in \Gamma : A(K) \subset U\}$ for all $K \subset X$ compact and $U \subset X$ open. The sets $B_{K,U}$ do not form a topological basis, but rather a sub-basis for the topology, so open sets in the compact-open topology are an arbitrary union of finite intersections of the sets $B_{K,U}$. A topological group Γ is called *discrete* if all points are open, and we say Γ acts *properly discontinuously* on X if for each compact subset $K \subset X$, the set $K \cap gK$ is non-empty for only finitely many $g \in \Gamma$. In the case where $\Gamma < PSL(2,\mathbb{R})$ acts on the hyperbolic plane, Γ is discrete if and only if Γ acts properly discontinuously. A group Γ acts *freely* on X if and only if each non-trivial $g \in \Gamma$ does not fix any points in X. A group $\Gamma < PSL(2,\mathbb{R})$ acts freely on D if and only if it is torsion-free [**Rat06**].

Let Γ be a freely acting, properly discontinuous group of isometries of the Klein model D. Then the quotient space D/Γ is the set of all Γ-orbits $\{\Gamma \cdot x : x \in D\}$ with the quotient metric $d_\Gamma(\Gamma \cdot x, \Gamma \cdot y) = \inf\{d(a,b) : a \in \Gamma \cdot x, b \in \Gamma \cdot y\}$. The quotient space is locally isometric to D and D is the universal covering space for D/Γ. The space $M = D/\Gamma$ is a hyperbolic surface with fundamental group $\pi_1(M) \cong \Gamma$, moreover, every hyperbolic surface arises this way [**Rat06**].

A subset R of a metric space X is a *fundamental domain* for the group Γ if R is connected and open, $g \cdot R : g \in \Gamma$ are pairwise disjoint and $X = \bigcup_{g\in\Gamma}\{g \cdot \bar{R} : g \in \Gamma\}$, where $\bar{R}$ is the closure of R in X.

1.8. Thick Thin Decomposition

Let M be a surface and $B(r)$ be a ball of radius r in the Euclidean plane. For a point x, the largest radius r such that the exponential map from $B(r)$ to the tangent space $T_x(M)$ is injective is called the *injectivity radius* of M at x. We denote the injectivity radius at x with $r(x)$. The injectivity radius $r(x)$ is also the supremum of radii of embedded metric balls in M centred at x.

We decompose M into a compact thick part

$$M^{\geq\varepsilon} = \{x \in M : r(x) \geq \varepsilon\}$$

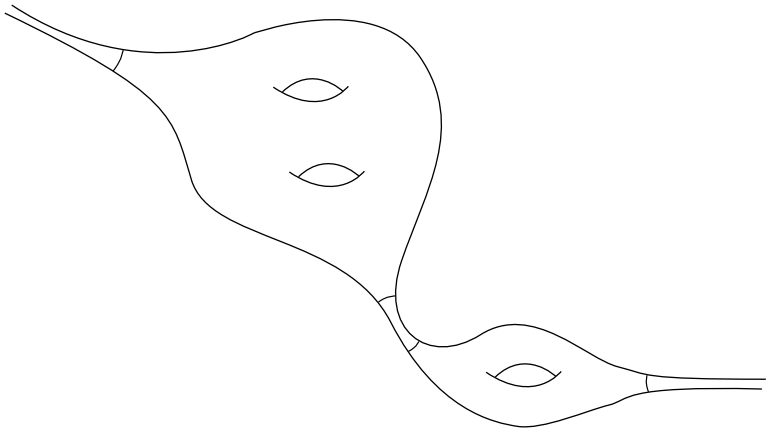

FIGURE 1.8.1. The thick thin decomposition of a surface.

and a thin part

$$M^{<\varepsilon} = \{x \in M : r(x) < \varepsilon\}.$$

To understand the thin part of this decomposition, we appeal to the following useful lemma.

Lemma 1.3 (Margulis Lemma)**.** *Let Γ be a group of discrete isometries of hyperbolic n-space. Then the group generated by*

$$\{\gamma \in \Gamma \mid d(x, \gamma(x)) \leq \varepsilon_n \ \ \forall x \in \mathbb{H}^n\}$$

is virtually nilpotent, where ε_n is a universal constant depending only on the dimension n.

Remark 1.4. An estimate for the Margulis constant $\varepsilon_2 \geq \operatorname{arcsinh}(1) \approx 0.8813$ is given in [**Bus10**].

Margulis lemma simplifies in the case of hyperbolic surfaces, as the manifold $\mathbb{H}^2/\Gamma$ has Γ torsion free so virtually nilpotent implies abelian. The abelian, discrete group is infinite cyclic with either a hyperbolic or parabolic generator. If we choose $\varepsilon < \varepsilon_2$, each component $M^{<\varepsilon}$ is an open annulus; either a regular neighbourhood of a short geodesic, or isometric to a subsurface of a surface of revolution, the *tractricoid* [**Thu97**]. This is called the *thick-thin decomposition* of surface M.

Explicitly, the tractricoid is parametrised by

$$t \mapsto (t - \tanh t, \operatorname{sech} t), \quad 0 \leq t < \infty$$

and is endowed with the metric d, defined by the property that $d(x, y)$ is equal to the infimum of the euclidean lengths of all piecewise differentiable curves joining x and y on the tractricoid. The metric d gives the tractricoid a complete hyperbolic structure with constant negative curvature.

1.9. Peripheral subgroup

If a non-compact surface S is topologically finite, then $S = \mathrm{Int}(\bar{S})$, where $\bar{S}$ is compact and the boundary $\partial\bar{S}$ is nonempty. For each component $\bar{c}_i \subset \partial\bar{S}$ take a neighbourhood $\bar{c}_i \times I \hookrightarrow \bar{S}$, called the *collar neighbourhood* of c_i, such that these neighbourhoods are pairwise disjoint. Denote $N(\partial\bar{S})$ the union of these neighbourhoods. Then $S^c := \bar{S} \backslash N(\partial\bar{S})$ is compact and homeomorphic to $\bar{S}$. This is called the *compact core* of S.

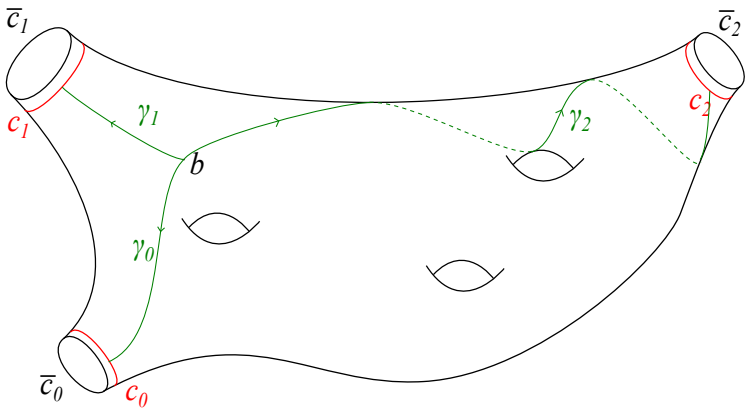

FIGURE 1.9.1. The collar neighbourhoods used to define the peripheral subgroups of $\pi_1(S, b)$.

Fix a basepoint b of S in S^c and choose a path $\gamma_i : [0, 1] \to S^c$ from b to c_i for each boundary component, where $\gamma_i(0) = b, \gamma_i(1) \in c_i$. Then there are natural homomorphisms $\phi_i : \pi_1(c_i, \gamma_i(1)) \to \pi_1(S^c, b) \cong \pi_1(S, b)$ by taking a loop α based at $\gamma_i(1)$ in the circle c_i to the loop $\gamma_i^{-1}\alpha\gamma_i$ in S based at b. Then $g^{-1}\,\mathrm{Im}(\phi_i)g$ is a *peripheral subgroup* of $\pi_1(S, b)$ for all $c_i \subset \partial S^c$ and $g \in \pi_1(S, b)$. This is well defined since choosing a different path γ_i results in a conjugate subgroup.

1.10. Parabolic fixed points

Consider a hyperbolic surface D/Γ, where Γ is a freely acting, discrete group of isometries. A point $p \in \partial D$ is a *parabolic fixed point* of Γ if there is a parabolic transformation $A \in \Gamma$ such that $Ap = p$. There is a one-to-one correspondence between the set of parabolic fixed points $\{p\}$ and the set of peripheral subgroups of the fundamental group Γ, as each peripheral subgroup fixes a unique point on the boundary. Equivalently, every parabolic fixed point of Γ is the fixed point of some peripheral subgroup of Γ. Consider a cusp of D/Γ with collar neighbourhood c_i. Its associated set of conjugate peripheral subgroups is defined to be $g^{-1}\,\mathrm{Im}\,\phi_i g$, where

$\operatorname{Im}\phi_i = \langle\gamma\rangle$ for a parabolic element $\gamma \in \Gamma$ if c_i is in the Margulis region $M^{<\varepsilon}$. If $\langle\gamma\rangle$ fixes the point $p \in \partial D$, then its conjugate $g^{-1}\langle\gamma\rangle g = \langle g^{-1}\gamma g\rangle$ fixes the point $g^{-1}p \in \partial D$ for all $g \in \Gamma$. Hence, each cusp of D/Γ has an associated set of conjugate peripheral subgroups with parabolic fixed points in a Γ-orbit.

For $p \in \partial D$, the *stabiliser subgroup* of p are elements of Γ which fix it, i.e.

$$\operatorname{Stab}_\Gamma(p) = \{A \in \Gamma : A \cdot p = p\}.$$

In the case where p is a parabolic fixed point of Γ, this stabiliser subgroup coincides with a peripheral subgroup of Γ.

1.11. Ideal cell decompositions

A *cell decomposition* of a manifold S with boundary ∂S is a decomposition of S as a union of cells (spaces homeomorphic $\mathbb{R}^n$, with any two cells meeting only along their boundaries).

An *ideal polyhedral decomposition* is a cell decomposition where every cell is an ideal polyhedron (polyhedron with its vertices removed). If the manifold is a non-compact surface and all 2-cells are ideal 2-simplices, the cell decomposition is called an *ideal triangulation* of the surface.

The following two lemmas are due to Lackenby [**Lac00**]:

Lemma 1.5. *Let S be a compact surface with non-empy boundary and $\chi(S) < 0$. Then $S \setminus \partial S$ admits an ideal triangulation with $-2\chi(S)$ ideal triangles.*

Lemma 1.6. *Any two ideal triangulations of a $S \setminus \partial S$ as above are related by a finite sequence of elementary moves. An elementary move consists of picking two distinct 2-simplices sharing an edge, removing the shared edge to form a square, and dividing this square along its other diagonal.*

1.12. The once punctured torus

The once punctured torus is the surface obtained by removing a point from the torus. It admits a Euclidean metric if considered as a restriction of the Euclidean torus, however, this metric is not complete. Although there is no complete Euclidean metric on the once punctured torus, it does admit a complete hyperbolic metric. The following construction can be found in [**Bon09**].

A cell decomposition of the once punctured torus is obtained by removing the vertices of a polyhedral decomposition of the torus. Hence, a ideal polyhedral decomposition of the once punctured torus is an ideal quadrilateral X with its opposite edges glued in an orientation preserving manner.

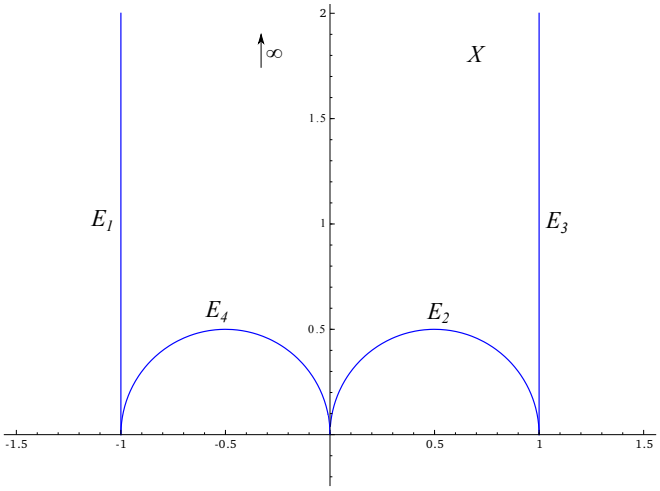

FIGURE 1.12.1. The ideal quadrilateral X with opposite edges glued is an ideal polyhedral decomposition of the once punctured torus

To give the once punctured torus a complete hyperbolic structure, X is embedded in the hyperbolic plane.

Let X be an ideal quadrilateral in the upper half-plane model $\mathbb{H}^2$, with ideal vertices at $-1, 0, 1, \infty$. The sides of X are the geodesics E_1 joining -1 to ∞, E_2 joining 0 to 1, E_3 joining 1 to ∞ and E_4 joining 0 to -1. Note that there is a one-parameter space of Möbius transformations which send the geodesic E_1 to E_2. The Möbius transformation

$$\phi_1(z) = \frac{z+1}{z+2}$$

sends the endpoints $-1, \infty$ to $0, 1$ respectively, so it sends E_1 to E_2. Similarly,

$$\phi_3(z) = \frac{z-1}{-z+2}$$

sends E_3 to E_4.

If $\phi_2 = \phi_1^{-1}$ and $\phi_4 = \phi_3^{-1}$, then the matrix forms of ϕ_4 and ϕ_2 in $PSL(2,\mathbb{R})$ are

$$A = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}, B = \begin{pmatrix} 2 & -1 \\ -1 & 1 \end{pmatrix}.$$

The two generator subgroup $\Gamma = \langle A, B \rangle$ of $PSL(2,\mathbb{R})$ is discrete and free in $PSL(2,\mathbb{R})$, and X is a fundamental domain for the group Γ. The quotient space $\widetilde{X} = \mathbb{H}^2/\Gamma$ is a hyperbolic surface homeomorphic to the once punctured torus. The commutator $[A, B]$ is $\phi_4 \cdot \phi_2 \cdot \phi_3 \cdot \phi_1$ and sends ∞ to

itself. Moreover,

$$[A, B] = \begin{pmatrix} -1 & -6 \\ 0 & -1 \end{pmatrix}$$

so $[A, B]$ is a parabolic element with its parabolic fixed point at ∞. The peripheral subgroups corresponding to the single cusp of the once punctured torus are generated by the conjugates of the commutator, i.e. $\langle g[A, B]g^{-1}\rangle$, where $g \in \Gamma$. The parabolic fixed points of Γ is the orbit $\Gamma \cdot \infty$.

Since X is convex, the quotient space $\widetilde{X}$ inherits a quotient metric which is hyperbolic. The thick thin decomposition of $\widetilde{X}$ consists of the compact thick part, and a thin part with a single component, being the neighbourhood of the single cusp, which can be described in terms of horoballs. Let $l_a = \{x + iy : y \geq a\}$ in $\mathbb{H}^2$ be a horoball where $a \geq 6/\varepsilon_2$. Then the peripheral subgroup of Γ fixing ∞ is $\langle [A, B]\rangle$, and acts discretely on the horoball l_a. The subset $Y = l_a/\Gamma \subset \mathbb{H}^2/\Gamma$ is a region in $\widetilde{X}$ in the neighbourhood of its cusp. Moreover, the injectivity radius of Y is the length of the boundary $\partial Y \subset X$. The peripheral element $[A, B]$ corresponds to the fractional linear transformation $z \mapsto z + 6$, so the length of the horocycle ∂l_a in the quotient space is $6 \cdot \frac{1}{a}$ as the metric on the horocycle is a scalar multiple of the Euclidean metric. Hence, the length of the boundary ∂Y is $\frac{6}{a} \leq \varepsilon_2$. The injectivity radius of any point in Y is less that ε_2, in fact, Y is the set of all such points. This gives the thick thin decomposition $M^{<\varepsilon} = Y$, $M^{>\varepsilon} = \widetilde{X}\backslash Y$.

The hyperbolic metric is complete in $\mathbb{H}^2$, so it is complete in the compact thick part $M^{>\varepsilon}$. The thin part is isometric to a subsurface of the tractricoid so the entire surface is hyperbolic and complete. [**Bon09**] gives an alternate elementary proof that the hyperbolic metric on Y is isometric to the natural metric on the tractricoid.

1.13. Developing map and holonomy

Let X be a connected real analytic manifold and G a group of real analytic diffeomorphisms acting transitively on X. A (G, X) manifold is a manifold M with a *geometric structure* given by a covering of M with charts $\phi_i : U_i \to X$ which map the open connected subsets U_i diffeomorphically onto open subsets of X. Whenever U_i and U_j overlap, the *transition map* is defined to be $\phi_j\phi_i^{-1} : \phi_i(U_i \cap U_j) \to \phi_j(U_i \cap U_j)$, and is the restriction of an element in G. In the case where G acts by isometries on X, this means that M is locally isometric to X.

If the geometric structure on (G, X) is analytic, then the analytic continuation of coordinate charts gives a global coordinate chart, or *developing map*, on the universal cover $\widetilde{M}$ of M. Given a basepoint x_0 and an initial

chart ϕ_0, the developing map is a map $D : \widetilde{M} \to X$ that agrees with the analytic continuation of ϕ_0 along each path and in the neighbourhood of the endpoints of the path. In other words, if $p : \widetilde{M} \to M$ is the universal covering map, $\alpha : [0, 1] \to M$ is a path with $\alpha(0) = x_0$, and ϕ_0^α is the analytic continuation of ϕ_0 along α, then the developing map $D = \phi_0^\alpha \cdot p$ along α and in a neighbourhood of $\alpha(1)$.

Let $\gamma \in \pi_1(M)$ and consider the initial chart ϕ_0 and the analytic continuation ϕ_0^γ along a loop representing γ. Since ϕ_0 and ϕ_0^γ are comparable as they are both defined in a neighbourhood of the basepoint x_0 in M. Since G acts transitively on X, there exists $g_\gamma \in G$ such that $\phi_0^\gamma = g_\gamma \phi_0$ when restricted to a neighbourhood of the base point. The map $\rho : \pi_1(X) \to G, \gamma \mapsto g_\gamma$ is a homomorphism and is called the *holonomy representation* for the geometric structure. A detailed discussion of the above is given in [**Thu97**].

CHAPTER 2

The Epstein-Penner Decomposition

2.1. Introduction

A decomposition of a manifold into cells, or a *cell decomposition*, is a useful tool in studying cusped hyperbolic n-manifolds. A particularly powerful construction is the Epstein-Penner convex hull construction [**EP$^+$88**], as it gives rise to canonical cell decompositions of finite volume, hyperbolic n-manifolds with $k \geq 1$ cusps. Applications of the Epstein-Penner construction include a cell decomposition of Teichmüller space [**Pen87**] and the creation of a database of cusped hyperbolic 3-manifolds [**CHW99**].

2.2. Convex hull construction

Suppose a hyperbolic surface S is connected, complete, topologically finite, has finite volume and has $k \geq 1$ cusps. Let $S \cong D/\Gamma$, where Γ is a torsion-free, discrete group of isometries, and assume Γ is finitely generated.

We lift D/Γ into Minkowski space by projecting its universal cover D to the hyperboloid model H. Then $S \cong D/\Gamma \cong H/\widetilde{\Gamma}$, where $\widetilde{\Gamma} < SO^+(1,2)$ is the group Γ conjugated by the radial projection map $\rho : D \to H$. We call this group isomorphism $\varphi : \Gamma \to \widetilde{\Gamma}, x \mapsto \rho x \rho^{-1}$. Parabolics $g \in \Gamma$ fixing the parabolic fixed point $p_i \in \partial D$ are lifted to a parabolic element $\varphi(g) \in \widetilde{\Gamma}$ fixing the ray $\lambda p_i \in L$, $\lambda \in \mathbb{R}$. In the case where Γ is represented by matrices in $PSL(2,\mathbb{R})$, the isomorphism $\varphi : \Gamma \to \widetilde{\Gamma}$ is the one given in section 1.6.

The cell decomposition of D/Γ proceeds as follows. A vector $\lambda_i p_i \in L^+$ chosen on the positive ray through p_i is called a *light-cone representative* of p_i. Choose a light-cone representative v_{p_i} for each parabolic fixed point p_i of Γ such that the set of light-cone representatives $B = \{v_{p_i}\}$ is $\widetilde{\Gamma}$-invariant. One way to ensure that B is $\widetilde{\Gamma}$-invariant is to choose a light-cone representative λp_i for a parabolic fixed point p_i, and then take the $\widetilde{\Gamma}$-orbit of λp_i to be the light-cone representatives for the parabolic fixed points $\Gamma \cdot p$. This ensures that B consists of $\widetilde{\Gamma}$-orbits, with one $\widetilde{\Gamma}$-orbit for each cusp of S.

Then by Theorem 2.2, B is discrete. Take the convex hull C of B and project the faces of ∂C radially onto the Klein model D. Since the construction of C is equivariant under the group $\widetilde{\Gamma}$, the faces of C are equivariant under $\widetilde{\Gamma}$, so the projection onto D is Γ-invariant and induces a cell decomposition of D/Γ.

Lemma 2.1. *There exists $\lambda \in \mathbb{R}$ so that the horoballs corresponding to λB are pairwise disjoint. In other words, the horoballs $\{h_w : w \in \lambda B\}$ are pairwise disjoint, where $h_w = \{x \in H : \langle w, x\rangle = -1\}$.*

Proof. If we take a thick thin decomposition $\{S^{\geq\varepsilon}, S^{<\varepsilon}\}$ of our surface S, where $\varepsilon < \varepsilon_2 \approx 0.88$, then each cusp of S lies in a different component of the thin part $S^{<\varepsilon}$. Suppose there exists $\lambda \in \mathbb{R}$ such that all horoballs $\{h_w : w \in \lambda B\}$ are in the thin part. In other words, $h_w/\Gamma \subset S^{<\varepsilon}$, i.e. injectivity radius $r(x)$ of $x \in h_w/\Gamma$ is less than ε.

The Γ-orbits of any horoball $h_w \subset S^{<\varepsilon}$ must be disjoint, since $x \in h_w \cap g \cdot h_w$, $g \in \Gamma$ implies there is a non-parabolic loop of length $< \varepsilon$ based at x, contradicting Margulis Lemma. For two horoballs h_{w_i} and h_{w_j} lying in distinct components, then h_{w_i} and h_{w_j} must be disjoint and their lift into the universal cover D will also be disjoint. Hence, if all horoballs $\{h_w : w \in \lambda B\}$ lie in the thin part, then the h_w are pairwise disjoint.

It remains to show that λ may be chosen such that all horoballs $\{h_w : w \in \lambda B\}$ are in the thin part. Equivalently, for each $x \in h_w/\Gamma$, the injectivity radius $r(x) < \varepsilon$. Hence, we need to calculate the injectivity radii in the space h_w/Γ, which can be conveniently done in the upper halfspace model $\mathbb{H}^2$.

Consider peripheral subgroups $\langle g^{-1}\gamma_i g\rangle$, $\gamma_i, g \in \Gamma$ of the i^{th}cusp. Now represent $\gamma_i \in \text{Isom}(\mathbb{H}^2)$ as an element in $PSL(2,\mathbb{R})$ and conjugate γ_i so that its parabolic fixed point is at ∞. Then $\gamma_i \in PSL(2,\mathbb{R})$ is $\left(\begin{smallmatrix} 1 & a \\ 0 & 1 \end{smallmatrix}\right)$. Note that this matrix in $PSL(2,\mathbb{R})$ is independent of which peripheral subgroup we chose for the i^{th} cusp since all peripheral subgroups are conjugate. The fractional linear transformation corresponding to $\left(\begin{smallmatrix} 1 & a \\ 0 & 1 \end{smallmatrix}\right)$ is the map $\tau_a = z \mapsto z + a$. Suppose we have a horoball based at ∞ quotiented by τ_a. Then the horoball is $l_b = \{x + iy : y > b\}$ modulo τ_a. The injectivity radius of $l_b/\tau_a \subset \mathbb{H}^2/\tau_a$ is maximal on its boundary. The metric on the horoball $\{x + ib\}$ is Euclidean with a scaling factor of $\frac{1}{b}$, hence, the length of the boundary of the horoball l_b/τ_a is $\frac{a}{b}$. Hence, $\frac{a}{b} < \varepsilon$, then the injectivity radius of points in l_b/τ_a will be less than ε.

Now we want to find the injectivity radii of $h_w/\Gamma \subset S$. Suppose w has first coordinate, or height, w_1 in $\mathbb{R}^3$. Then by Lemma 1.2, this horocycle has Euclidean radius $\frac{1}{1+w_1}$ in the Poincaré Disk model. Furthermore, by Lemma 1.1, if the horocycle is conjugated to be centred at ∞, it would be $l_{w_1} = \{x + iy : y > w_1\}$ in the upper halfspace model. But $w_1 > \frac{a}{\varepsilon}$

implies that l_{w_1}/τ_a has injectivity radius $r(x) < \varepsilon$. Hence, for h_w/Γ to have injectivity radius less than ε, we need the first coordinate w_1 to be greater than $\frac{a}{\varepsilon}$.

Hence, define λ as follows. For each $\widetilde{\Gamma}$-orbit $B_i \subset B$ choose a point $w \in B_i$. The choice of $w \in B_i$ does not matter as they all correspond to the same horoball S. Calculate the peripheral element which fixes w, and find its conjugate representation $\left(\begin{smallmatrix} 1 & a_i \\ 0 & 1 \end{smallmatrix}\right)$ in $PSL(2,\mathbb{R})$. Then choose $\lambda_i > \frac{a_i}{\varepsilon w_1}$, $\lambda_i \in \mathbb{R}$. The height of λw is $\lambda_1 w_1 > \frac{a_i}{\varepsilon}$, so the horoball corresponding to λw has injectivity radius $< \varepsilon$. Hence, the horoball corresponding to B_i lies in the Margulis region $S^{<\varepsilon}$. Hence if we define $\lambda = \max_{i=1}^{k} \lambda_i$, then the horoball corresponding to the orbit B_i will be in the thin part of S. Hence, all horoballs $h_w = \{x \in H : \langle w, x\rangle = -1\}$ are pairwise disjoint. □

Theorem 2.2. *The set B is discrete.*

Proof. We will instead prove that λB is discrete for the $\lambda \in \mathbb{R}$ provided by Lemma 2.1. Suppose λB is not discrete, hence there is an accumultation point. There is a lower halfspace $T_y = \{x = (x_1, x_2, x_3)\} \in \mathbb{R}^3, x_1 < y$ which includes the point of accumulation, so $\lambda B \cap T_y$ has infinitely many points.

Any point $w \in \lambda B \cap T_y$ has height $w_1 < y$. Let w be the light-like representative of the horoball $h_w,$. Then by Lemma 1.2, the horoball h_w has Euclidean radius $\frac{1}{1+w_1}$ in the Poincaré Disk model. So for every $w \in \lambda B \cap T_y$, its corresponding horoball h_w has Euclidean radius at least $\frac{1}{1+y}$ in the Poincaré model. But the horoballs h_w with light-like representatives in λB are disjoint, by Lemma 2.1. Hence, we have infinitely many disjoint horoballs of Euclidean radius at least $\frac{1}{1+y}$, and this is impossible since the Poincaré Disk has finite area. □

Theorem 2.3. *The convex hull construction associates a $(k-1)$-parameter family of Γ-invariant cell decompositions of D. In particular if $k = 1$ then the decomposition of the surface D/Γ is canonical.*

Proof. The cell decomposition of D/Γ can be altered if we change the $\widetilde{\Gamma}$ orbits B_i to $\alpha_i B_i$, where $\alpha_i \in \mathbb{R}^+, 1 \leq i \leq k$. However, if $\alpha_1 = \alpha_2 = ... = \alpha_k$, the orbit B is dilated by α_1 so its projection onto D is unchanged. Therefore we have a $(k-1)$-parameter family of cell decompositions of D/Γ. □

Remark 2.4. If $k \geq 2$ and we choose $\alpha_i B_i$ so that their corresponding horoballs bound the same area in S, then we have a canonical cell decomposition for S.

2.3. Example

Take the once punctured torus with complete hyperbolic structure defined in Section 1.12. The space is given by $\mathbb{H}^2/\Gamma$ where $\Gamma = \langle A, B\rangle$,

$$A = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}, B = \begin{pmatrix} 2 & -1 \\ -1 & 1 \end{pmatrix},$$

with commutator

$$ABA^{-1}B^{-1} = \begin{pmatrix} -1 & -6 \\ 0 & -1 \end{pmatrix}.$$

so $[A, B]$ is parabolic with parabolic fixed point at ∞.

Now we let $\mathbb{H}^2/\Gamma \cong H/\widetilde{\Gamma}$ where H is the hyperboloid model and $\widetilde{\Gamma} < SO^+(1,2)$ is a discrete, freely acting subgroup with two generators. The construction of $\widetilde{\Gamma}$ from $\Gamma \subset PSL(2,\mathbb{R})$ is given in 1.6, and the matrices corresponding to A, B are

$$\widetilde{A} = \begin{pmatrix} \frac{7}{2} & 3 & \frac{3}{2} \\ 3 & 3 & 1 \\ \frac{3}{2} & 1 & \frac{3}{2} \end{pmatrix}, \widetilde{B} = \begin{pmatrix} \frac{7}{2} & -3 & \frac{3}{2} \\ -3 & 3 & -1 \\ \frac{3}{2} & -1 & \frac{3}{2} \end{pmatrix}.$$

The commutator is

$$[\widetilde{A}, \widetilde{B}] = \begin{pmatrix} 19 & 6 & 18 \\ 6 & 1 & 6 \\ -18 & -6 & -17 \end{pmatrix},$$

and since

$$\begin{pmatrix} 19 & 6 & 18 \\ 6 & 1 & 6 \\ -18 & -6 & -17 \end{pmatrix} = P \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix} P^{-1}, \quad P = \begin{pmatrix} 36 & 18 & 1 \\ 0 & 6 & 0 \\ -36 & -18 & 0 \end{pmatrix}$$

the commutator is indeed a parabolic element of $SO^+(1,2)$. The one dimensional invariant subspace is $\lambda(1,0,-1)$, $\lambda \in \mathbb{R}$, so the commutator $[\widetilde{A}, \widetilde{B}]$ is a peripheral element of $\widetilde{\Gamma}$ with a fixed point at $v_p = (1,0,-1)$. The peripheral subgroups of $\widetilde{\Gamma}$ are conjugate to $\langle[\widetilde{A}, \widetilde{B}]\rangle$. As there is only one cusp in the once punctured torus, these are the only peripheral subgroups. Hence, set of the parabolic fixed points of $\widetilde{\Gamma}$ is the $\widetilde{\Gamma}$ orbit of v_p. To obtain the canonical cell decomposition, take the convex hull of this orbit and projecting it onto D/Γ.

However, calculating the convex hull of the infinite orbit $\widetilde{\Gamma} \cdot v_p$ is computationally infeasible if done naively. We will need the edge flipping algorithm in Chapter 5 or Weeks' Algorithm [**Wee93**] to calculate the canonical cell decomposition efficiently.

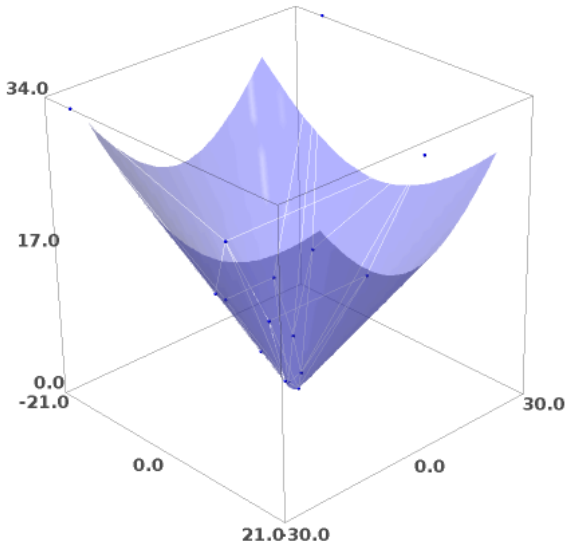

FIGURE 2.3.1. The orbit $N^3 v_p$ on the light-cone (blue).

Without such algorithms, the best we can do is to approximate the infinite convex hull by computing finitely many points in the orbit. Let $N = \{Id, \widetilde{A}, \widetilde{B}, \widetilde{A}^{-1}, \widetilde{B}^{-1}\}$ be a neighbourhood of the identity in the Cayley graph of $\widetilde{\Gamma}$ with respect to the natural generating set. Then N^k consists of the group elements distance at most k from the identity. Hence, $N^k v_p$ are parabolic fixed points on the light-cone (see Figure 2.3.1) and its convex hull (see Figure 2.3.2) gives a finite approximation to the convex hull of Γv_p. Projecting the convex hull of $N^3 v_p$ onto the Klein model gives an approximation to the the canonical cell decomposition (see Figure 2.3.3).

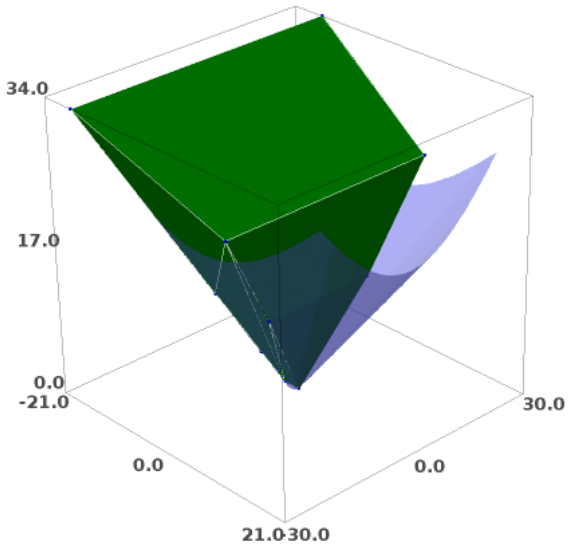

FIGURE 2.3.2. The orbit $N^3 v_p$ and its convex hull (green).

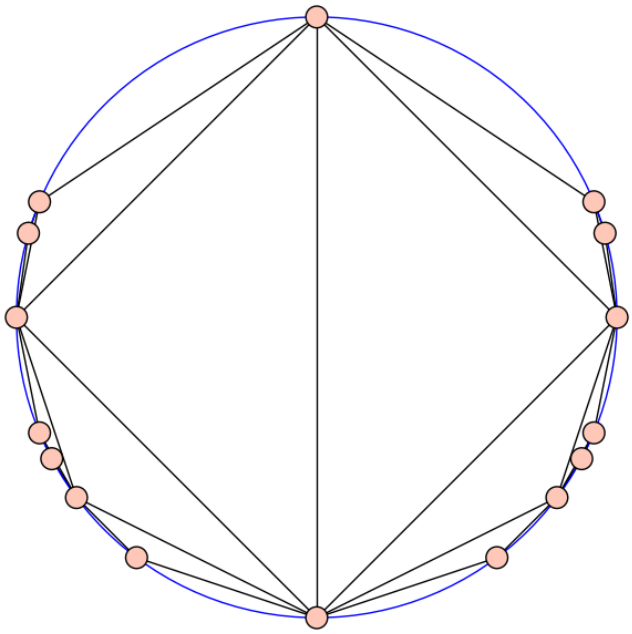

FIGURE 2.3.3. The projection of the convex hull of $N^3 v_\infty$ onto the Klein model D.

CHAPTER 3

Background - Projective Geometry

3.1. The real projective plane

Let $\sim$ be an equivalence relation on $x, y \in \mathbb{R}^3$, where $x \sim y$ if and only if $x = \lambda y$ for some $\lambda \in \mathbb{R}\backslash\{0\}$. Then points are equivalent if and only if they lie on the same line through the origin. The quotient

$$\mathbb{R}P^2 := (\mathbb{R}^3\backslash\{0\})/\sim$$

is called the *real projective plane*. Therefore, the real projective plane $\mathbb{R}P^2$ is the set of all lines in $\mathbb{R}^3$ passing through the origin, and lines in $\mathbb{R}P^2$ correspond to planes in $\mathbb{R}^3$.

Each $\sim$-equivalence class of $\mathbb{R}P^2$ intersects the unit sphere $S^2 \subset \mathbb{R}^3$ at two points. Therefore, another definition of the real projective plane is $\mathbb{R}P^2 := S^2/\pm 1$ with quotient map $\pi : S^n \to \mathbb{R}P^2$. The quotient map is a double cover in that for any $a \in \mathbb{R}P^2$, there is a unique pair of antipodal points such that $\pi(x) = a$ and $\pi(-x) = a$.

Any linear transformation of $\mathbb{R}^3$ preserves equivalence classes of $\sim$ so it induces a tranformation of $\mathbb{R}P^2$. The group of transformations of $\mathbb{R}P^2$, called *projective transformations*, is the *projective general linear group* $PGL(3, \mathbb{R})$. The group $PGL(3, \mathbb{R})$ is isomorphic to $GL(3, \mathbb{R})$ quotiented by scalar multiplication. Furthermore, the map $A \mapsto (\det A)^{\frac{1}{3}} A$ defines an isomorphism between $PGL(3, \mathbb{R})$ and $SL(3, \mathbb{R})$, where $SL(3, \mathbb{R})$ is the group of 3×3 matrices of determinant 1.

Let $A, B, C, D \in \mathbb{R}P^2$ be collinear, in other words, the lines $A, B, C, D \in \mathbb{R}^3$ are coplanar. Let P be a line in R^3 which intersects A, B, C, D at four distinct points a, b, c, d in that order. Then the *cross ratio* of A, B, C, D is defined to be

$$(A, B; C, D) = \frac{|ac||bd|}{|ab||cd|},$$

where $|\cdot|$ is the Euclidean distance. The cross ratio is a projective invariant in that it is independent of P and is preserved by any projective transformation.

3.2. Strictly convex domains

A set $C \in \mathbb{R}P^2$ is called *convex* if the intersection of any projective line with C is connected. An *affine patch* of is a subset of $\mathbb{R}P^2$ by deleting a projective line. A convex set is *properly convex* if its closure is contained in an affine patch. A properly convex set is called *strictly convex* if its boundary does not contain any lines of positive length.

For any strictly convex domain $\Omega \subset \mathbb{R}P^2$, $\pi^{-1}(\Omega)$ consists of two copies of X in S^2. If Ω lies in an affine patch $A \subset \mathbb{R}P^2$, then $\pi^{-1}(\Omega) \subset \pi^{-1}(A) = Hem^+ \cup Hem^-$, where $Hem^+, Hem^- \subset S^2$ are two open hemispheres with boundary $\partial Hem^+ = \partial Hem^- = l$ where $l \subset S^2$ is a great circle. Hence, we have two disjoint components of $\pi^{-1}(\Omega)$, we identify Ω with the component $\Omega^+ \in Hem^+$ and define the *positive light-cone* of Ω to be $\mathcal{L} = \mathbb{R}^+ \cdot \partial\Omega^+ \in \mathbb{R}^3$.

Any strictly convex set Ω has a complete metric defined as follows: identify Ω with a component $\Omega^+ \subset Hem^+$ in the double cover and project Ω^+ radially onto a plane P that does not pass through the origin. Let the radial projection of points $p, q \in \Omega^+$ and domain Ω^+onto P be $, \bar{p}, \bar{q}$ and the domain $\bar{\Omega}$. Let the line $\bar{p}\bar{q}$ intersect the boundary $\partial\bar{\Omega}$ at points $\bar{a}$, $\bar{b}$ so that $\bar{a}, \bar{p}, \bar{q}, \bar{b}$ are in that order, and define the distance between $p, q \in \Omega$ be defined by the cross ratio

$$d(p, q) = \frac{1}{2} \log \frac{|\bar{q}\bar{a}||\bar{b}\bar{p}|}{|\bar{p}\bar{a}||\bar{q}\bar{b}|},$$

where $|\cdot|$ denotes the Euclidean distance on P. Note that the disance $d(p, q)$ is independent of choice of projection plane P as cross ratios are projectively invariant. The metric d is called the *Hilbert Metric* on the strictly convex set Ω.

Note that the metric on the Klein model coincides with its Hilbert metric, hence, the hyperbolic geometry coincides with geometry on a projective disk. A strictly convex domain Ω has hyperbolic geometry if and only if its projection $\bar{\Omega}$ is isometric to the Klein disk, i.e. when Ω is an ellipse.

3.3. Projective Transformations of a strictly convex domain

The group of projective transformations on the strictly convex domain Ω are the elements of $PGL(3, \mathbb{R})$ which fix Ω, which we will denote by $PGL(\Omega)$.

Then $PGL(\Omega)$ is isomorphic to group of linear transformations of $\mathbb{R}^3$ which stabilise the component $\Omega^+ \in \mathbb{R}P^2$. Hence, $PGL(\Omega) \cong SL_{\pm}(\Omega)$, where $SL_{\pm}(\Omega)$ is the group of 3×3 matrices of determinant ± 1 which stabilises $\Omega^+ \subset \mathbb{R}^3$.

Similarly to the hyperbolic setting, we define an projective transformation $A \in PGL(\Omega)$ to be *elliptic* if A fixes a point in Ω. Otherwise, A acts freely on Ω and is called *parabolic* if every eigenvalue has modulus 1 and *hyperbolic* otherwise. A point $p \in \partial\Omega$ is called a *parabolic fixed point* of a group Γ if there is a parabolic element $A \in \Gamma$ such that $Ap = p$.

3.4. Convex Projective Manifolds

Any strictly convex, finite volume real projective surface can be written as $S = \Omega/\Gamma$, where Ω is strictly convex and $\Gamma \cong \pi_1(S) \subset PGL(\Omega)$ is a free, discrete group of projective transformations. The Hilbert metric defines a metric on S and the notion of volume.

A thorough discussion of strictly convex projective manifolds and cusps, as well as proofs for Theorems 3.1 and 3.2, can be found in [**CLT11**].

Theorem 3.1 simplifies for projective surfaces. Virtually nilpotent mplies the subgroup of $\pi_1(S, x)$ is infinite cycle with either a hyperbolic of parabolic generator. Respectively, these correspond to the thin component of x being an annular neighbourhood of a short geodesic or cusp end. If S has finite volume, then the cusp end is projectively equivalent to hyperbolic cusps, meaning their holonomies are conjugate to $\left(\begin{smallmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{smallmatrix}\right)$.

Theorem 3.1. *(Properly Convex Margulis) For $n \geq 2$ there is a Margulis constant μ_n such that any properly convex projective n-manifold M and any point $x \in M$, the subgroup of $\pi_1(M, x)$ generated by loops based at x of length less than μ_n is virtually nilpotent.*

Theorem 3.2. *(Strictly Convex Thick-Thin) Let M be a finite volume, strictly convex projective n-manifold. Then $M = A \cup B$, where A and B are smooth submanifolds and $\bar{A} \cap \bar{B} = \partial A = \partial B$, B is compact and non-empty and A is possibly empty with the following properties:*

a) *If $\mathrm{inj}(x) < \nu_n$ then $x \in A$, where $\mathrm{inj}(x)$ is the injectivity radius of M at x, and $\nu_n = 3^{-(n+1)}\mu_n$.*
b) *If $x \in A$ then $\mathrm{inj}(x) < \mu_n$.*
c) *Each component of A is a Margulis tube or a cusp.*

3.5. The projective once punctured torus

Theorem 3.3. *There is a 6 dimensional parameter space of projective structures on the once punctured torus, given by the points*

$$Q = \{(c_1, c_2, b_1, a, b, e) \in \mathbb{R}^6 | c_1 > 1,\ c_2 > 1,\ a_2 b_1 > 1,\ a, b, e > 0\}$$

where

$$a_2 = \frac{\begin{array}{c} a^3b^3e^6 + a^3bb_1e^5 - a^6b_1c_1c_2 - 2\,a^3b^2e^4 - b^2e^7 + a^6b_1c_1 \\ + a^3bc_1c_2e^2 + \big(a^6b^2c_1 - a^6b^2 - \big(a^6b^2c_1 - a^6b^2\big)c_2\big)e \end{array}}{\begin{array}{c} a^3b^3b_1e^6 - b^4e^8 + \big(a^3b^3c_1 - a^3b^3\big)c_2e^3 \\ + \big(a^6b^2b_1c_1 - a^6b^2b_1 - \big(a^6b^2b_1c_1 - a^6b^2b_1\big)c_2\big)e \end{array}}.$$

The ideal polyhedral decomposition of these once punctured tori is an ideal quadrilateral with opposite edges glued. Fixing the four vertices of the ideal quadrilateral

$$\begin{pmatrix}1\\0\\0\end{pmatrix}, \frac{1}{\sqrt{3}}\begin{pmatrix}1\\1\\-1\end{pmatrix}, \begin{pmatrix}0\\1\\0\end{pmatrix}, \begin{pmatrix}0\\0\\1\end{pmatrix} \in S^2.$$

fixes the gluing matrices up to conjugation. The point (c_1, c_2, b_1, a, b, e) *corresponds to the pair of matrices*

$$\widetilde{E} = \begin{pmatrix} ab_1 + \frac{ac_1}{be^2} & a & \frac{a}{be^2} \\ ab_1 - \frac{be^2}{aa} & a & 0 \\ -ab_1 & -a & 0 \end{pmatrix},$$

$$\widetilde{F} = \begin{pmatrix} a_2b & 0 & a_2b - \frac{1}{be} \\ -b & 0 & -b \\ c_2b & e & bc_2 + e \end{pmatrix}.$$

which represent orientation preserving isometries with parabolic commutator. Furthermore, for each point $(c_1, c_2, b_1, a, b, e) \in Q$*, there is a strictly convex domain* $\Omega \in \mathbb{R}P^2$ *such that* $\Omega^+/\langle \widetilde{E}, \widetilde{F}\rangle$ *homeomorphic to the once punctured torus.*

Proof. Let T be a once punctured torus with fundamental group $\pi_1(T) = \langle E, F\rangle$ and let $\rho : \pi_1(T) \to SL_\pm(3, \mathbb{R})$ be a representation of the fundamental group into $SL_\pm(3, \mathbb{R})$. Assume further that there exists a strictly convex domain $\Omega \subseteq \mathbb{R}P^2$ such that ρ corresponds to a strictly convex projective structure, i.e. $\rho(\pi_1(T)) < SL_\pm(\Omega)$ and $T \cong \Omega/\rho(\pi_1(T))$.

Let a fundamental domain of T be given by the convex quadrilateral $Y = \{p_1, p_2, p_3, p_4\}$ in the double cover S^2 of the projective plane, with the points in clockwise order and corresponding to the puncture.

Then, $\rho(E), \rho(F)$ are the gluing equations of Y which make $\Omega/\langle E, F\rangle$ the projective once punctured torus.

Now let $h \in SL_\pm(3, \mathbb{R})$ be a matrix satisfying

$$h(p_1) = e_1 = \begin{pmatrix}1\\0\\0\end{pmatrix}, h(p_2) = e_2 = \begin{pmatrix}0\\1\\0\end{pmatrix}, h(p_3) = e_3 = \begin{pmatrix}0\\0\\1\end{pmatrix}.$$

Then the conjugate representation $h\rho h^{-1}$ has parabolic fixed points at those standard basis vectors. Hence, if we consider only representations up to conjugation we can assume without loss of generality that the vertices of Y are $\{e_1, p, e_2, e_3\}$ for some $p \in S^2$.

Consider $g \in GL(3, \mathbb{R})$ of the form

$$g = \begin{pmatrix} \lambda & 0 & 0 \\ 0 & \mu & 0 \\ 0 & 0 & \nu \end{pmatrix}$$

where $\lambda, \mu, \nu \in \mathbb{R}^+$. Then g fixes the rays through e_1, e_2, e_3, and g sends p to any element in S^2 we want. Hence we can assume that the fourth vertex of Y is

$$p = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix}.$$

This fixes the vertices $\{e_1, p, e_2, e_3\}$ of convex quadrilateral $Y \subset S^2$ and fixes ρ in its conjugacy class.

Now the pair of gluing equations on Y are $E, F \in \pi_1(T)$, where E sends $e_2 \to p$, $e_3 \to e_1$ and F sends $p \to e_1$, $e_2 \to e_3$. Then in terms of $SL_{\pm}(3, \mathbb{R})$ we have

$$\widetilde{E} \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} = a \begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix}, \qquad \widetilde{E} \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} = d \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$$

and

$$\widetilde{F} \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} = e \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \qquad \widetilde{F} \begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix} = m \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$$

where $\widetilde{E} = \rho(E)$, $\widetilde{F} = \rho(F)$ and $a, d, e, m \in \mathbb{R}^+$. To fix $\widetilde{E}, \widetilde{F} \in SL_{\pm}(3, \mathbb{R})$ we need one more equation for each of E and F.

Let

$$\widetilde{E} \begin{pmatrix} u \\ v \\ w \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \qquad \widetilde{F} \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} x \\ y \\ z \end{pmatrix}$$

where $x, y, z, u, v, w \in \mathbb{R}$.

Since the quadrilateral with vertices

$$\begin{pmatrix} u \\ v \\ w \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$$

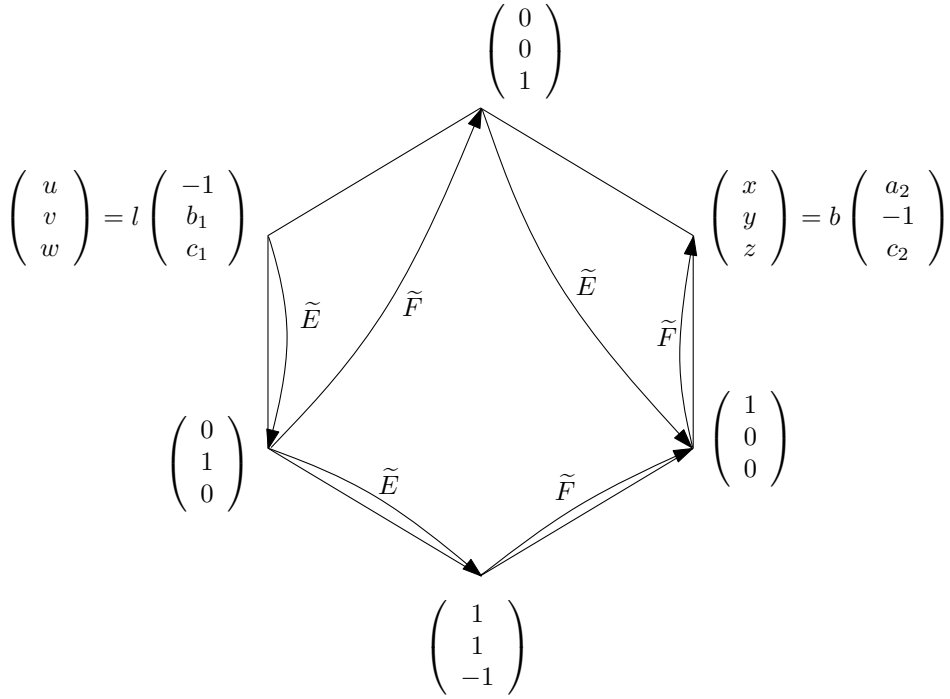

FIGURE 3.5.1. The action of $\widetilde{E}$, $\widetilde{F}$ on the ideal vertices of Y and two additional points, forming a convex hexagon.

is convex, the points $\left(\begin{smallmatrix} u \\ v \\ w \end{smallmatrix}\right)$, $\left(\begin{smallmatrix} 0 \\ 1 \\ 0 \end{smallmatrix}\right)$ must lie on the same side of the great circle through $\left(\begin{smallmatrix} 0 \\ 0 \\ 1 \end{smallmatrix}\right)$, $\left(\begin{smallmatrix} 1 \\ 0 \\ 0 \end{smallmatrix}\right)$. Hence, $u > 0$. Similarly, we obtain the inequalities $u < 0$, $v > 0$, $w > 0$, $x > 0$, $y < 0$, $z > 0$. We reparametrise our variables

$$(u, v, w) = (-l, lb_1, lc_1), \qquad (x, y, z) = (a_2 b, -b, c_2 b),$$

where $l, b_1, c_1, b, a_2, c_2 \in \mathbb{R}^+$. Furthermore, cross-ratio invariants give the inequalities $c_2 > 1, c_1 > 1, a_2 b_1 > 1$, and details can be found in [**G$^+$90**]. The set-up is slightly different but turns out to be equivalent. Now we can write the matrix form of $\widetilde{E}$ and $\widetilde{F}$.

$$\widetilde{E} = \begin{pmatrix} a & d & 0 \\ a & 0 & 1 \\ -a & 0 & 0 \end{pmatrix} \begin{pmatrix} 0 & 0 & -l \\ 1 & 0 & lb_1 \\ 0 & 1 & lc_1 \end{pmatrix}^{-1} = \begin{pmatrix} ab_1 + c_1 d & a & d \\ ab_1 - \frac{1}{l} & a & 0 \\ -ab_1 & -a & 0 \end{pmatrix}$$

$$\widetilde{F} = \begin{pmatrix} 0 & m & a_2 b \\ 0 & 0 & -b \\ e & 0 & c_2 b \end{pmatrix} \begin{pmatrix} 0 & 1 & 1 \\ 1 & 1 & 0 \\ 0 & -1 & 0 \end{pmatrix}^{-1} = \begin{pmatrix} a_2 b & 0 & a_2 b - m \\ -b & 0 & -b \\ c_2 b & e & bc_2 + e \end{pmatrix}$$

Since $\widetilde{E}, \widetilde{F} \in SL(3,\mathbb{R})$, we have two equations from $\det(\widetilde{E}) = \frac{ad}{l}$ and $\det(\widetilde{F}) = bem$. Hence we can let $l = ad$ and $m = \frac{1}{be}$ to give

$$\widetilde{E} = \begin{pmatrix} ab_1 + c_1 d & a & d \\ ab_1 - \frac{1}{ad} & a & 0 \\ -ab_1 & -a & 0 \end{pmatrix},$$

$$\widetilde{F} = \begin{pmatrix} a_2 b & 0 & a_2 b - \frac{1}{be} \\ -b & 0 & -b \\ c_2 b & e & bc_2 + e \end{pmatrix}.$$

This gives a parametrisation of all pairs of matrices gluing the opposite of Y together in S^2. We also require the ideal vertices of Y on $\partial\Omega$ to be parabolic fixed points, so the peripheral subgroup corresponding to e_1 must be a parabolic transformation. The peripheral element at e_1, by inspection, is $EFE^{-1}F^{-1}$, and the representation in $SL_{\pm}(3,\mathbb{R})$ is of the form

$$[\widetilde{E}, \widetilde{F}] = \begin{pmatrix} \frac{bde^2}{a} & * & * \\ 0 & * & * \\ 0 & * & * \end{pmatrix}.$$

where $*$ are algebraic expressions in terms of our 8 remaining variables. For the commutator to be parabolic with a parabolic fixed point at e_1, we must have $[\widetilde{E}, \widetilde{F}]e_1 = e_1$ and $\mathrm{tr}[\widetilde{E}, \widetilde{F}] = 3$. The first condition enforces $d = \frac{a}{be^2}$ whereas the second conditions gives an equation linear in a_2. Hence we can eliminate a_2 with the reparametrisation

$$a_2 = \frac{\begin{array}{c} a^3b^3e^6 + a^3bb_1e^5 - a^6b_1c_1c_2 - 2\,a^3b^2e^4 - b^2e^7 + a^6b_1c_1 \\ + a^3bc_1c_2e^2 + \big(a^6b^2c_1 - a^6b^2 - \big(a^6b^2c_1 - a^6b^2\big)c_2\big)e \end{array}}{\begin{array}{c} a^3b^3b_1e^6 - b^4e^8 + \big(a^3b^3c_1 - a^3b^3\big)c_2e^3 \\ + \big(a^6b^2b_1c_1 - a^6b^2b_1 - \big(a^6b^2b_1c_1 - a^6b^2b_1\big)c_2\big)e \end{array}}$$

Moreover, with this substitution, all eigenvalues of $[\widetilde{E}, \widetilde{F}]$ are 1 and the Jordan canonical form of the commutator is the standard parabolic

$$[\widetilde{E}, \widetilde{F}] \sim \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}.$$

Hence, we have a 6 parameter space of projective structures, and the representation $\widetilde{E}, \widetilde{F}$ always has parabolic commutator. Moreover, all peripheral subgroups of $\langle \widetilde{E}, \widetilde{F} \rangle$ are of the form $\langle g[\widetilde{E}, \widetilde{F}]g^{-1} \rangle$ where $g \in \langle \widetilde{E}, \widetilde{F} \rangle$.

Now, by Corollary 4.6 in [**Mar10**], since our representation is discrete and faithful, peripheral subgroups are parabolic and the space is homeomorphic to the once punctured torus, the space Q is homeomorphic to the moduli space of marked properly convex projective structures on a finite volume once punctured torus. □

CHAPTER 4

Generalisation of the Epstein-Penner Construction to Projective Manifolds

4.1. Cooper-Long's Construction

The Cooper-Long Construction [**CL13**] is a generalisation of the Epstein-Penner construction in the case of strictly convex projective manifolds. Its construction is analogous to the Epstein-Penner construction in Chapter 2.

Let Ω be a strictly convex domain and suppose $S = \Omega/\Gamma$ is a finite volume, strictly convex projective surface with $k \geq 1$ cusps. For $p \in \partial\Omega$, call any point $\lambda p \in \mathcal{L} = \mathbb{R}^+ \cdot \partial\Omega^+$ a *light-cone representative* of p.

For each parabolic fixed point $p_i \in \partial\Omega$ of Γ, take a light-cone representative v_{p_i} of p_i so that the set $B = \{v_{p_i}\}$ is Γ-invariant. This can be done by choosing a light-cone representative λp_i for a parabolic fixed point p_i, and taking its Γ-orbit in $\mathcal{L}$. There is one such Γ-orbit for each of the k cusps of S.

Since B is discrete, let C be convex hull of B. The polyhedron C is Γ-invariant since B is Γ-invariant. Therefore, the projection of the faces ∂C onto Ω is Γ-invariant and induces a cell decomposition of Ω/Γ.

Theorem 4.1. *B is discrete.*

Proof. See [**CL13**]. The argument is a generalisation of the argument in Theorem 2.2. The difficulty is that horospheres are defined via Busemann functions and are not circles in general, making the volume argument less straightforward. □

Theorem 4.2. *The convex hull construction associates a $(k-1)$-parameter family of Γ-invariant cell decompositions of Ω. In particular if $k = 1$ then the decomposition of the surface Ω/Γ is canonical.*

Proof. Same as Theorem 2.3. □

4.2. Example

Consider the projective structure corresponding to the point $(8, 2, 3, 7, 1, 2)$ in Theorem 3.3. The corresponding matrices can be calculated to be:

$$\widetilde{E} = \begin{pmatrix} 35 & 7 & \frac{7}{4} \\ \frac{1025}{49} & 7 & 0 \\ -21 & -7 & 0 \end{pmatrix}, \qquad \widetilde{F} = \begin{pmatrix} \frac{2202467}{2418621} & 0 & \frac{1986313}{4837242} \\ -1 & 0 & -1 \\ 2 & 2 & 4 \end{pmatrix}.$$

The commutator

$$[\widetilde{E}, \widetilde{F}] = \begin{pmatrix} 1 & \frac{11437407629}{9674484} & \frac{11508050281}{19348968} \\ 0 & \frac{5871445}{806207} & \frac{19811401}{4837242} \\ 0 & -\frac{7770264}{806207} & -\frac{4259031}{806207} \end{pmatrix}$$

is indeed parabolic, since

$$[\widetilde{E}, \widetilde{F}] = P \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix} P^{-1}, \qquad P = \begin{pmatrix} 1 & \frac{11437407629}{9674484} & \frac{11508050281}{19348968} \\ 0 & \frac{5871445}{806207} & \frac{19811401}{4837242} \\ 0 & -\frac{7770264}{806207} & -\frac{4259031}{806207} \end{pmatrix}.$$

Hence, by Theorem 3.3, the space Ω/Γ is a projective once punctured torus, where $\Gamma = \langle \widetilde{E}, \widetilde{F} \rangle$ and Ω is some strictly convex subset of $\mathbb{R}P^2$.

The commutator $[\widetilde{E}, \widetilde{F}]$ has a unique one dimensional invariant subspace $(\lambda, 0, 0)$. Since the once punctured torus has only one cusp, the Cooper-Long convex hull construction is just the convex hull of the Γ-orbit of any light-cone representative, e.g. $e_1 = (1, 0, 0)$.

The convex hull C cannot be calculated naively as there are infinitely many vertices. We must appeal to the edge flipping algorithm in Chapter 5 to compute the canonical cell decomposition in finite time. Note that Weeks' Algorithm in [**Wee93**] does not apply in this case since Weeks' tilt formula is specific to the hyperbolic setting.

Again, we will do the best we can without the algorithm in Chapter 5, and approximate the convex hull C by computing finitely many points on the orbit. Let $N = \{Id, \widetilde{E}, \widetilde{F}, \widetilde{E}^{-1}, \widetilde{F}^{-1}\}$ and take the finite orbit $N^k e_1$ (see Figure 4.2.1). Then take the convex hull C_k of the orbit $N^k e_1$ (see Figure 4.2.2). Project the faces of C_k onto the sphere S^2. These faces are on an affine patch Hem^+. We flatten the affine patch by taking the orthogonal projection of points on Hem^+ onto the plane through ∂Hem^+ (see Figure 4.2.3).

Notice in Figure 4.2.3 we have finitely many points on the boundary Ω^+, and this gives us an approximate shape for the domain Ω^+. The parabolic fixed points of Γ is dense in Ω^+ for any Γ discrete and freely acting. So with enough points in our orbit $B \in \mathcal{L}$, we can obtain an arbitrarily fine approximation to the boundary of Ω^+.

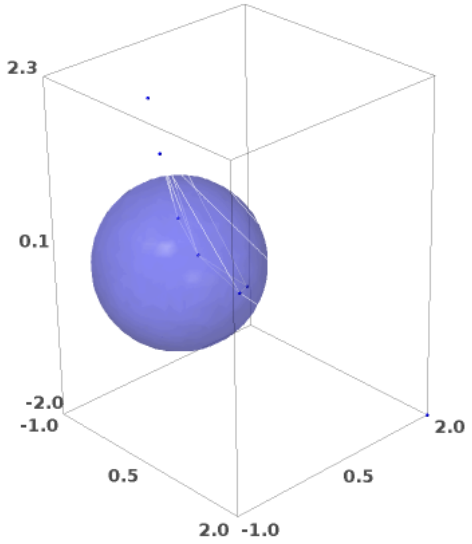

FIGURE 4.2.1. The orbit $N^3 e_1$.

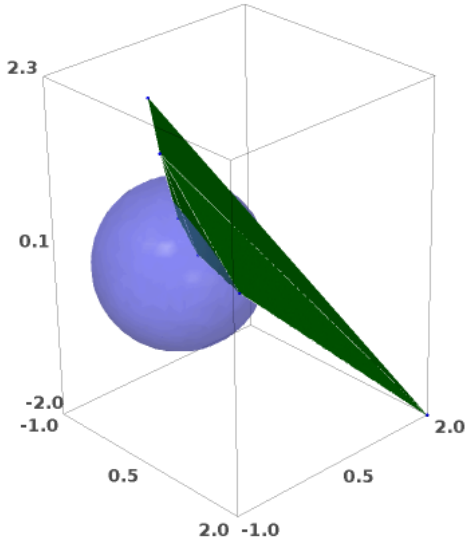

FIGURE 4.2.2. The orbit $N^3 e_1$ and its convex hull (green).

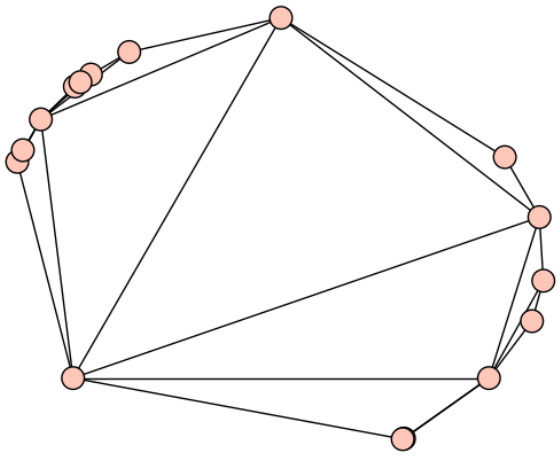

FIGURE 4.2.3. The projection of the convex hull of $N^3 v_\infty$ onto an affine patch of S^2.

CHAPTER 5

Edge Flipping Algorithm for the Epstein-Penner Construction

5.1. Introduction

The first use of an edge flipping algorithm was by Lawson [**Law77**] to compute Delauney triangulations. Edge flipping algorithms use a sequence of local modifications to arrive at a globally optimal solution, and decisions on which edge to flip come from purely local considerations. Edge flipping algorithms also work well in computing convex hulls [**GCTH13**], however, these algorithms are not applicable as our case involve infinitely many vertices.

Weeks' algorithm [**Wee93**] is an edge flipping algorithm to compute the canonical cell decomposition of a cusped n-manifold. Weeks' algorithm is only proven to be correct for $n = 2$, in higher dimensions it is only a heuristic procedure. The algorithm is included in the software package SnapPea which has many applications, such as a census of all cusped finite volume hyperbolic 3-manifolds [**Bur14**] or computing arithmetic invariants of hyperbolic 3-manifolds [**CGHN00**].

In this chapter an alternative algorithm is provided, its advantages over Weeks' algorithm are it is simpler to calculate, and may be extended to strictly convex projective surfaces.

5.2. The Algorithm

The crucial difference between the proposed algorithm and Weeks' is the method of detection of convex angles. The Weeks' tilt formula [**Wee93**] and its generalised versions [**SW95**], [**Ush02**] rely heavily on the Minkowski metric whereas the proposed algorithm does not use the underlying metric of the space.

The new method is to checks the following property of a convex hull: if F is a face of the convex hull C, all other vertices of the polyhedron are on the same side of X_F, where X_F is the plane through F. In the case where C is the convex hull of the Γ-invariant discrete subset B on a light-cone $\mathcal{L}$, the plane X_F separates $\mathbb{R}^3$ into two components. If F is a face of C, then

the interior of the polyhedron C lies entirely in one component of $\mathbb{R}^3 \backslash X_F$, and the other component contains the origin.

Define the *neighbouring faces* of F to be the faces which share an edge with F. Define the vertex of a neighbouring face which is not on the shared edge to be a *neighbouring vertex* of F. Let X_F be the plane passing through F. If v and the origin lie in the same component of $\mathbb{R}^3 \backslash X_F$, then we say v is *below* the face F. If v and the origin lie in different components then v is *above* the face F. Otherwise v is on the plane X_F and is *coplanar* with F.

Define an *edge flip* on F and a neighbouring vertex v to be the edge flip which removes the common edge. If the vertices of F are $\{a, b, c\}$ where ab is the common edge, then the edge flip creates the two new faces $\{c, v, a\}$ and $\{c, b, v\}$. We call an edge flip *admissible* if v is below F.

We call F *locally convex* if each neighbouring vertex v is either above F or coplanar with F. Equivalently, F is locally convex if there are no admissible edge flips which include face F.

Let Ω be a strictly convex domain in $\mathbb{R}P^2$, for example, the Klein model of hyperbolic space. The edge flipping algorithm is as follows. Let the projective surface be Ω/Γ where $\Gamma \subset SL_{\pm}(\Omega)$ is freely acting, discrete and finitely generated. We start with an arbitrary cell decomposition of Ω/Γ into ideal triangles, its existence is ensured by Lemma 1.5. The ideal triangulation projects to a Γ-invariant polyhedron with Γ-invariant vertices on the light-cone $\mathcal{L}$.

For a face F on the Γ-invariant polyhedron, we call $\Gamma \cdot F$ the *face class* of F. Proposition 5.1 shows that F is locally convex if and only if gF is convex, where $g \in \Gamma$. Hence, it makes sense to call a face class $\Gamma \cdot F$ locally convex.

For each face class $\Gamma \cdot F$, we check the neighbouring vertices for any admissible edge flips. If there is an admissible edge flip, then it is performed (replacing $\Gamma \cdot F$ and another face class with two different face classes). This gives a different Γ-invariant polyhedron with vertices on the light-cone $\mathcal{L}$, and the entire algorithm starts again.

If there are no admissible edge flips, another face class is checked. Although there are infinitely many faces in the polyhedron, there are only finitely many face classes.

The algorithm terminates when there are no more admissible edge flips. The algorithm terminates in finitely many steps (Theorem 5.5). Moreover, the Γ-invariant polyhedron in the final iteration is convex (Proposition 5.3).

There is one final step in the algorithm to make the polyhedron equal to C. Even though our polyhedron is convex, since each face class is a triangle, the polyhedron is actually a *triangulation* of ∂C. We call an edge of a polyhedron *redundant* if it lies in the interior of a face of ∂C. Note that an edge is redundant if and only if the two faces that meet at e are coplanar.

After all admissible edge flips are performed, a cleanup step is applied. If two adjacent faces are coplanar, then their common edge is removed. This is repeated until there are no redundant edges, upon which the polyhedron is equal to the convex hull (Theorem 5.4).

Pseudocode for this algorithm is provided below.

Algorithm 1 Edge Flipping Algorithm for the Convex Hull Construction

```
1: T = an arbitrary cell decomposition of Ω/Γ into triangles
2: P = Γ-invariant polyhedron induced by T.
3: for all F ∈ P/Γ̃ do
4:    for all v a neighbouring vertex of F do
5:       if v is below F then
6:          Perform an edge flip on {v, F}
7:          go to line 3
8:       end if
9:    end for
10: end for
11: for all e an edge of P/Γ̃ do
12:    if F_1, F_2 sharing e are coplanar then
13:       Remove e and merge F_1, F_2
14:    end if
15: end for
16: return P
```

5.3. Proof of correctness

Proposition 5.1. *Point $\widetilde{p}$ lies below face $F \iff g\widetilde{p}$ lies below gF for all $g \in SL_{\pm}(\Omega)$.*

Proof. Let the face F have vertices $\widetilde{f_1}, \widetilde{f_2}, \widetilde{f_3}$ on the light-cone $\mathcal{L}$, and let their projections onto the boundary $\partial\Omega$ be f_1, f_2, f_3 respectively. Let the projection of $\widetilde{p}$ onto $\partial\Omega$ be p and without loss of generality assume p, f_1, f_2, f_3 are in clockwise order around $\partial\Omega$.

Let the segments pf_2 and f_1f_3 intersect at point x. Then $x \in \Omega$ since Ω is strictly convex. Since p, x, f_2 are collinear, the rays through p, x, f_2 are coplanar, so segment $\widetilde{p}\widetilde{f_2}$ passes through a point λx, $\lambda \in \mathbb{R}^+$. Similarly, segment $\widetilde{f_1}\widetilde{f_3}$ passes through a point μx, $\mu \in \mathbb{R}^+$, but $\mu > \lambda$ since this is the only case where $\widetilde{p}$ is below face F. In particular this means that the segment $\widetilde{p}\widetilde{f_2}$ intersects the interior of triangle $o\widetilde{f_1}\widetilde{f_3}$, where o is the origin (see Figure 5.3.1).

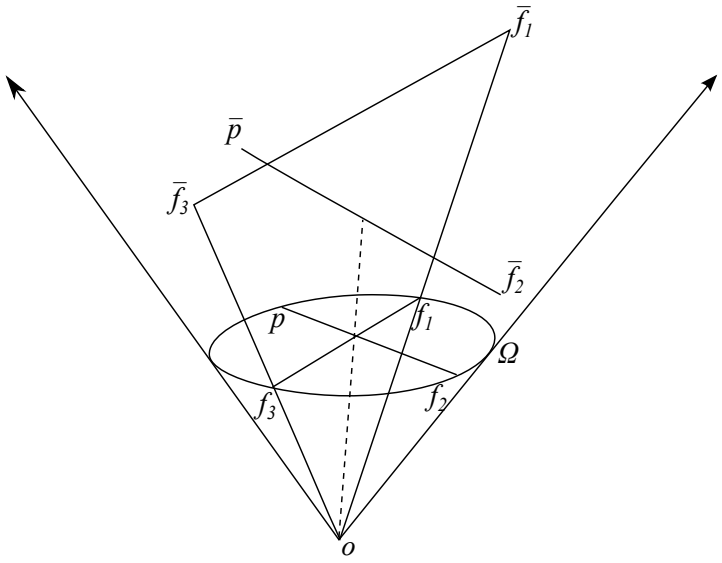

FIGURE 5.3.1.

Hence, p lying below face F is equivalent to the segment $\widetilde{p}\widetilde{f_2}$ intersecting the interior of triangle $o\widetilde{f_1}\widetilde{f_3}$ for some ordering $\{\widetilde{f_1}, \widetilde{f_2}, \widetilde{f_3}\}$ of the vertices of F.

However, the intersection point of the segment and the triangle can be written both as a convex combination of $\widetilde{p}$, $\widetilde{f_2}$ and a convex combination of $o, \widetilde{f_1}, \widetilde{f_3}$. This property is linear and preserved by any linear transformation $g \in GL_3(\mathbb{R})$.

Hence, if $g \in SL_{\pm}(\Omega)$, then g preserves the light-cone $\mathcal{L}$ and gF has its vertices on $\mathcal{L}$. Moreover, since the convex combination property is linear and preserved by g, $\widetilde{p}$ lies below face F if and only if $g\widetilde{p}$ lies below gF. □

Remark 5.2. Proposition 5.1 gives an alternative way to prove that the convex hull construction in Sections 2.2 and 4.1 are Γ-invariant. If F is a face on the convex hull, then there are no other vertices of the polyhedron below F, so the same is true for gF for all $g \in \Gamma$. Hence, F is a face of the convex hull if and only if gF is a face of the convex hull, so the convex hull is Γ-invariant.

Proposition 5.3. *If polyhedron P is locally convex at every face class, then the polyhedron is globally convex. In particular, if for every face F, v lies above face F for every neighbouring vertex v of F, then for every face F, u lies above face F for every vertex u of the entire polyhedron P.*

Proof. Suppose that polyhedron P is locally convex at every face, but is not globally convex. There exists a point $\widetilde{p}$ and a face $\widetilde{a}\widetilde{b}\widetilde{c}$ such that $\widetilde{p}$ lies

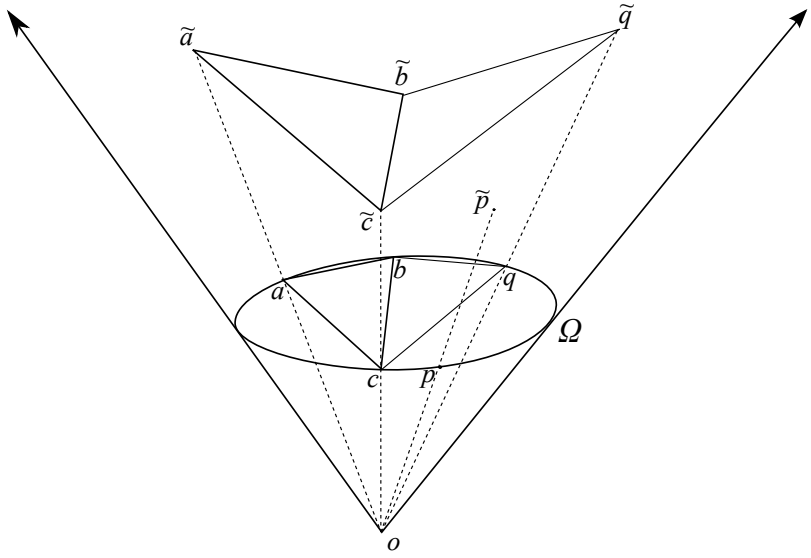

FIGURE 5.3.2.

below $\widetilde{a}\widetilde{b}\widetilde{c}$. Without loss of generality let the projection of $\widetilde{a}\widetilde{b}\widetilde{p}\widetilde{c}$ be a, b, p, c in clockwise order around $\partial\Omega$.

Let the face of P sharing edge $\widetilde{b}\widetilde{c}$ with F have vertices $\widetilde{b}, \widetilde{q}, \widetilde{c}$. Since $\widetilde{q}$ is a neighbouring vertex of $\widetilde{a}\widetilde{b}\widetilde{c}$, $\widetilde{q}$ lies above the hyperplane passing through $\widetilde{a}\widetilde{b}\widetilde{c}$. In particular, $\widetilde{p}\widetilde{q}$ lie on opposite sides of the faces $\widetilde{a}\widetilde{b}\widetilde{c}$ (see Figure 5.3.2).

Consider hyperplanes through $\widetilde{a}\widetilde{b}\widetilde{c}$ and $\widetilde{b}\widetilde{q}\widetilde{c}$, which intersect along the line $\widetilde{b}\widetilde{c}$. Let the hyperplane through $o\widetilde{b}\widetilde{c}$ divide $\mathbb{R}^3$ into two half-spaces, with T^+ including $\widetilde{p}, \widetilde{q}$ and T^- including $\widetilde{a}$. Then $\widetilde{a}\widetilde{b}\widetilde{c}$ lies above $\widetilde{b}\widetilde{q}\widetilde{c}$ in the half-space T^-, whereas $\widetilde{b}\widetilde{q}\widetilde{c}$ lies above $\widetilde{a}\widetilde{b}\widetilde{c}$ in the half-space T^+. Hence, $\widetilde{b}\widetilde{q}\widetilde{c}$ lies above $\widetilde{a}\widetilde{b}\widetilde{c}$ which lies above p in the half-space T^+, so $\widetilde{p}$ lies below $\widetilde{b}\widetilde{q}\widetilde{c}$.

Hence, we have $\widetilde{p}$ lying below $\widetilde{b}, \widetilde{q}, \widetilde{c}$ instead of $\widetilde{a}\widetilde{b}\widetilde{c}$, where $\widetilde{b}\widetilde{q}\widetilde{c}$ is "closer" to $\widetilde{p}$ than $\widetilde{a}\widetilde{b}\widetilde{c}$. Note that "closer" is well defined, as there is a unique path of triangles from $\widetilde{a}\widetilde{b}\widetilde{c}$ to $\widetilde{p}$ (this is clear if we project onto Ω). Also the path of triangles has finite length. Hence, if we initially assume that $\widetilde{p}$ is a point which lies below P, and $\widetilde{a}\widetilde{b}\widetilde{c}$ is the closest face to $\widetilde{p}$ such that $\widetilde{p}$ is below it, then we have another triangle with $\widetilde{p}$ below it, contradicting the minimality assumption. Hence, if P is locally convex at every face, then P is globally convex.

□

Theorem 5.4. *If Algorithm 1 terminates, the output P is the convex hull of the orbit B.*

Proof. For each face class $F \in P/\Gamma$, Algorithm 1 checks that the polyhedron is locally convex at F. So P is locally convex at every face, so by Proposition 2.3, P is convex. Hence the union of the faces of P is equal to ∂C. After the cleanup step, there will be no redundant edges and $P = C$. □

Theorem 5.5. *Algorithm 1 terminates in finitely many iterations.*

Proof. Let B be the Γ-invariant vertices of P. Define the height of a face F to be the number of points the orbit B below the face. The height of any face is finite since B is discrete, in particular 0 is not an accumulation point. Moreover, the height of a face is Γ-invariant by Proposition 5.1. Hence, we can define the height of the polyhedron P to be the sum of the heights of its face classes $\Gamma \cdot F \in P/\Gamma$.

To prove that Algorithm 1 terminates, it suffices to show that the height of P strictly decreases after every edge flip, since the height of P is always a non-negative integer.

Consider an edge flip $\widetilde{ac} \to \widetilde{bd}$. This occurs only if $\widetilde{c}$ lies above $\widetilde{abd}$ and $\widetilde{a}$ lies above $\widetilde{bcd}$. Without loss of generality let their projections onto $\partial\Omega$ be a, b, c, d in clockwise order. Let $\widetilde{p}$ be a point on the light-cone and without loss of generality let the projections p, a, b, c, d be in clockwise order.

Suppose $\widetilde{p}$ is below $\widetilde{abd}$. Then consider the triangles $\widetilde{abd}$ and $\widetilde{abc}$. The hyperplane through $o, \widetilde{a}, \widetilde{b}$ divides $\mathbb{R}^3$ into two half-spaces, consider only the halfspace which includes $\widetilde{c}, \widetilde{d}, \widetilde{p}$. In this halfspace, we know by local convexity that triangle $\widetilde{abc}$ is above $\widetilde{abd}$, which in turn is above $\widetilde{p}$. Hence, $\widetilde{p}$ is below $\widetilde{abd}$ implies that $\widetilde{p}$ is also below $\widetilde{abc}$ (see Figure 5.3.3).

Hence, when only considering points with projection between d, a in the strictly convex domain Ω, the points below $\widetilde{abd}$ is a subset of points below $\widetilde{abc}$. By exactly the same argument the points below $\widetilde{bcd}$ are a subset of points below $\widetilde{acd}$ when considering this particular part on the boundary $\partial\Omega$. An equivalent result is true the four other arcs ab, bc and cd by rotating the edge labels cyclically. However, the points $\widetilde{b}, \widetilde{d}$ contribute to the heights of $\widetilde{acd}, \widetilde{abc}$ before the edge flip, whereas no vertices contribute to the heights of the two triangles $\widetilde{abd}, \widetilde{bcd}$. Hence, after every edge flip, the height of the two triangles strictly decreases. The heights of all other triangles not involved in the edge flip stay the same, so the overall result is that the height of P strictly decreases after every edge flip. □

Remark 5.6. This edge flipping algorithm for the Epstein-Penner construction gives an alternative proof to Lemma 1.6.

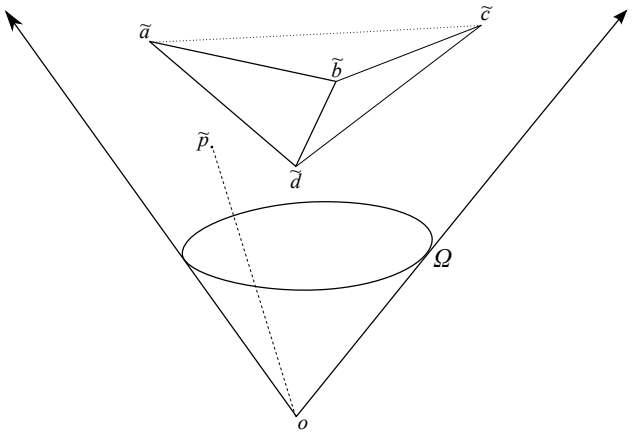

FIGURE 5.3.3.

If S is the interior of a compact surface with boundary and $\chi(S) < 0$, then it has a complete hyperbolic structure of finite volume. Then the Epstein-Penner construction provides a canonical cell decomposition of S. Moreover, Theorem 5.4 together with Theorem 5.5 implies that any ideal triangulation of S is can be modified into the canonical cell decomposition by a finite number of edge flips and deleting a finite number of redundant edges. Since deleting edges is not an elementary move, the two theorems imply that any ideal triangulation of S is related to the canonical cell decomposition but with its polygonal cells triangulated.

But any two triangulations of a polygon are also related by a finite number of edge flips. Hence, any two ideal triangulations are related to the canonical cell decomposition plus redundant edges, so are related to each other.

5.4. Examples

Summary. In this section a number of examples are examined:

- The once punctured torus in Section 1.12
- An extension of the previous example which requires multiple edge flips. By generalising this method, we may generate examples which require arbitrarily many edge flips.
- An example where the cell decomposition consists of only one cell, an ideal quadrilateral. This corresponds to a "critical point" in the moduli space, and neighbouring points have different cell decompositions.

- The projective once punctured torus in Section 4.2 corresponding to the point $(8, 2, 3, 7, 1, 2) \in Q$.
- Another projective once punctured torus, this time corresponding to the point $(4, 7, \frac{3}{2}, 1, 1, 7) \in Q$
- A one-parameter space of projective structures lying on a segment joining
$$(8, 2, 7, 2, 4, 2), (8, 2, 3, 7, 1, 2) \in Q$$

These examples are the inputs to the edge flipping algorithm. Software for this algorithm was implemented by the author in `sage` and took 906 lines of code.

Example 1. Take the once punctured torus given in Section 1.12. Our cusped manifold has a fundamental domain with ideal vertices at p, Ap, Bp, ABp in the Klein model, where $p = \infty$ is a parabolic fixed point. An initial cell decomposition of the fundamental domain consists of two triangles $\{p, Ap, Bp\}$ and $\{Ap, Bp, ABp\}$, and is the input to the edge flipping algorithm. This corresponds to a Γ-invariant polyhedron with two face classes $\{p, \widetilde{A}p, \widetilde{B}p\}$ and $\{\widetilde{A}p, \widetilde{B}p, \widetilde{A}\widetilde{B}p\}$. We will call this polyhedron P_1.

The algorithm detects that $\widetilde{A}\widetilde{B}v_p$ is below $\{v_p, \widetilde{A}v_p, \widetilde{B}v_p\}$. An edge flip on the triangle $\{\widetilde{p}, \widetilde{A}p, \widetilde{B}p\}$ and its neighbouring vertex $\widetilde{A}\widetilde{B}p$ is performed, giving a new triangulation. The new Γ-invariant polyhedron is called P_2.

The algorithm terminates after one edge flip as all remaining face classes are locally convex. P_2 is the convex hull and induces the canonical cell decomposition of D. The projection of P_1/Γ and P_2/Γ onto the Klein model D is shown in Figures 5.4.1 and 5.4.2.

These cell decompositions of D/Γ may be developed into cell decompositions of D, as shown in Figures 5.4.1 and 5.4.2.

Example 2. Take the same once punctured torus as in Example 1. Let P_3 be the polyhedron after applying an edge flip to a *non-admissable* edge. Instead of taking us to the canonical cell decomposition, this takes us further away from it in that there are more edge flips required to get to the convex hull. Continue applying non-admissible edge flips a number of times to give a Γ-invariant polyhedron P_n. Use P_n as the input to Algorithm 1.

P_n requires many edge flips for the algorithm the terminate, in fact P_n may need arbitrarily many edge flips. Consider a graph where points are the possible Γ-invariant polyhedron of the Γ-orbit of v_p, where $v_p = (1, 0, -1)$ and $\Gamma = \langle A, B\rangle$ defined in Section 1.12. Then there are infinitely many points in this graph, since there are infinitely many choices of a fundamental domain for D/Γ. Each polyhedron has 3 possible edge flips, so the degree of each vertex in the graph is 3. Hence, there are at most 4^n points of distance at most n from the convex hull, so there are infinitely many points

with distance $\geq n$ for any n. Hence, P_n may need arbitrarily many edge flips to arrive at the convex hull.

An example constructed using the method above is shown in Figure 5.4.5, and needs 3 edge flips to reach the convex hull. Since the Epstein-Penner decomposition is canonical, the final cell decomposition should match with the one in Example 1. This is difficult to notice by comparing Figure 5.4.2 and Figure 5.4.5, as the fundamental domains of the two figures do not match. However, if we develop the cell decomposition of D/Γ to a cell decomposition of D, then it is clear that the two cell decompositions do indeed match(see Figure 5.4.4 and Figure 5.4.6).

Example 3. In this example, the convex hull C has non-triangular faces, and hence the cleanup step is necessary. One possibility where C has a non-triangular face is if representations A, B can be found such that p, $\widetilde{A}p$, $\widetilde{B}p$ and $\widetilde{A}\widetilde{B}p$ are coplanar in $\mathbb{R}^3$, this results in a single quadrilateral cell in its canonical cell decomposition. Take the 2-parameter space of complete hyperbolic structures on the once punctured torus as given in [**Ser99**]:

$$A = \begin{pmatrix} \frac{z^2+1}{w} & z \\ z & w \end{pmatrix},\ B = \begin{pmatrix} \frac{w^2+1}{z} & -w \\ -w & z \end{pmatrix}.\ z, w \in \mathbb{R}$$

Then, the points p, $\widetilde{A}p$, $\widetilde{B}p$ and $\widetilde{A}\widetilde{B}p \in \mathbb{R}^3$ can be calculated in terms of z and w. Solving for the case where the four points are coplanar gives the relation

$$z = \sqrt{1 - w^2}.$$

The cell decomposition in Figure 5.4.8 corresponds to the parameters $w = 0.6$, $z = \sqrt{1 - w^2} = 0.8$. The resulting cell decomposition is a single ideal quadrilateral, as expected.

Figure 5.4.7 is the cell decomposition corresponding to the parameters $w = 0.6$, $z = 0.799$ whereas Figure 5.4.9 is corresponds to the $w = 0.6$, $z = 0.801$. This small perturbation in opposite direction results in two different cell decompositions. This is expected since the vertices of the quadrilateral $\{\widetilde{p}, \widetilde{A}p, \widetilde{B}p, \widetilde{A}\widetilde{B}p\}$ are in a degenerate position, so $w = 0.6$, $z = 0.8$ corresponds to a "critical point" of the parameter space.

Example 4. Take the projective once punctured torus in Section 4.2, corresponding to the point $(8, 2, 3, 7, 1, 2) \in Q$. We apply the edge flipping algorithm to find the canonical cell decomposition of $\Omega/\langle \widetilde{E}, \widetilde{F} \rangle$. An initial triangulation consists of the two triangles with vertices $\{e_1, \widetilde{E}e_1, \widetilde{F}e_1\}$ $\{\widetilde{E}e_1, \widetilde{F}e_1, \widetilde{E}\widetilde{F}e_1\}$. Two edge flips are performed, after which no more neighbouring vertices are below their respective faces, and our algorithm terminates. Figure 5.4.10 show the polyhedra visited by Algorithm 1, in particular, the figure shows its projection of its face classes onto an affine

patch of S^2. Figure 5.4.11 shows the cell decomposition of Ω/Γ developed into a cell decomposition of Ω. The figure shows more than one face for each face class. The ideal vertices of the triangulation are points on the boundary. Since the set of parabolic fixed points is dense on $\partial\Omega$, generating many points in the orbit B gives an approximation to the shape of Ω. In particular, by inspection, the boundary looks strictly convex and similar to an ellipse.

Example 5. Take the projective once punctured torus with structure given by $(4, 7, \frac{3}{2}, 1, 1, 7) \in Q$ in Theorem 3.3. Then its representing matrices in $SL_{\pm}(3, \mathbb{R})$ are

$$\widetilde{E} = \begin{pmatrix} \frac{29}{3} & 1 & \frac{2}{3} \\ \frac{11}{2} & 1 & 0 \\ -7 & -1 & 0 \end{pmatrix}, \qquad \widetilde{F} = \begin{pmatrix} \frac{85}{69} & 0 & \frac{13}{23} \\ -\frac{3}{2} & 0 & -\frac{3}{2} \\ \frac{21}{2} & 1 & \frac{23}{2} \end{pmatrix}$$

with commutator

$$[\widetilde{E}, \widetilde{F}] = \begin{pmatrix} 1 & \frac{182365}{1242} & \frac{13985}{621} \\ 0 & \frac{787}{207} & \frac{160}{207} \\ 0 & -\frac{4205}{414} & -\frac{373}{207} \end{pmatrix}$$

being a parabolic element of $SL_{\pm}(3, \mathbb{R})$ since

$$[\widetilde{E}, \widetilde{F}] = P \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix} P^{-1}, \qquad P = \begin{pmatrix} \frac{15654925}{85698} & \frac{182365}{1242} & 0 \\ 0 & \frac{580}{207} & 1 \\ 0 & -\frac{4205}{414} & 0 \end{pmatrix}.$$

We take the light-cone representative $e_1 = (1, 0, 0)$, which is a parabolic fixed point of the peripheral element $[\widetilde{E}, \widetilde{F}]$. The initial triangulation, and input to the edge flipping algorithm, is the pair of triangles $\{e_1, \widetilde{E}e_1, \widetilde{F}e_1\}$ and $\{\widetilde{E}e_1, \widetilde{F}e_1, \widetilde{E}\widetilde{F}e_1\}$. The algorithm terminates after one edge flip (see Figure 5.4.12). Figure 5.4.13 shows the cell decomposition of Figure 5.4.12 developed in the universal cover. Unlike the previous example, the boundary is not similar to an ellipse.

To check that the boundary is not an ellipse, we fit an ellipse through 5 points on the boundary, as shown in Figure 5.4.14 (note that it takes 5 points to define an ellipse since two ellipses may intersect at 4 distinct points). Clearly, the ellipse does not pass through all the points on the boundary. Recall that a projective structure is hyperbolic if and only if the boundary is the central projection of a circle, hence the projective structure defined by $(4, 7, \frac{3}{2}, 1, 1, 7) \in Q$ is not hyperbolic.

Example 6. If we take two projective structures, for example

$$(8, 2, 7, 2, 4, 2), (8, 2, 3, 7, 1, 2) \in Q,$$

one way to relate the two is to examine the structures in between. In particular, we can look at the cell decompositions of the linear combinations

$$(1 - \lambda)(8, 2, 7, 2, 4, 2) + \lambda(8, 2, 3, 7, 1, 2) \in Q.$$

This shows us how the cell decomposition of one slowly changes into the other.

Figure 5.4.15 shows the canonincal cell decompositions of 12 evenly spaced sample points along the line joining $(8, 2, 7, 2, 4, 2)$ and $(8, 2, 3, 7, 1, 2)$ in the six dimensional parameter space Q. Only the canonical cell decompositions are shown. By inspection we can deduce that an edge flip occured along this segment, in particular, between $\lambda = 0.82$ and $\lambda = 0.91$. A binary search with more sample points would improve this estimate further.

It is worth noting that attempts to solve for the parameter λ where the edge flip occured turned out to be equivalent to solving a degree 13 polynomial. Hence, approximating the position where the edge flip occurs is the best we can do.

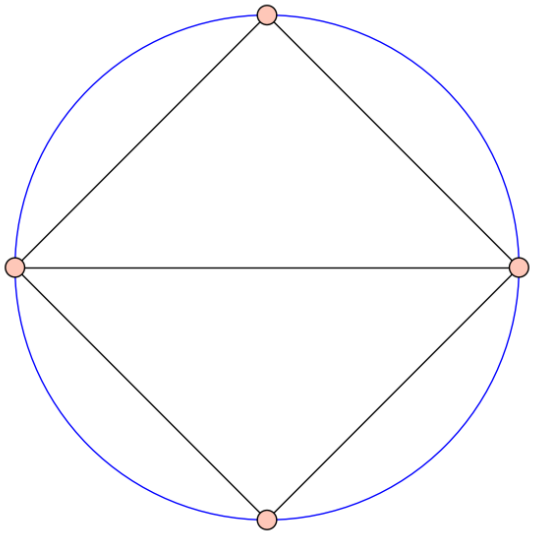

FIGURE 5.4.1. (Example 1) Projection P_1/Γ onto the Klein model.

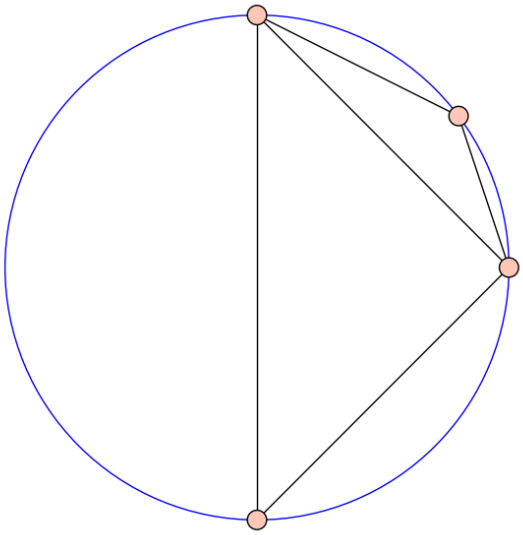

FIGURE 5.4.2. (Example 1) Projection of P_2/Γ onto the Klein model.

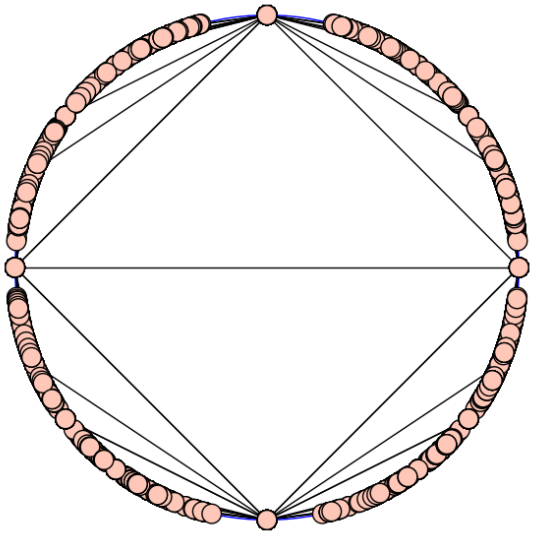

FIGURE 5.4.3. (Example 1) Projection of P_1 onto the Klein model

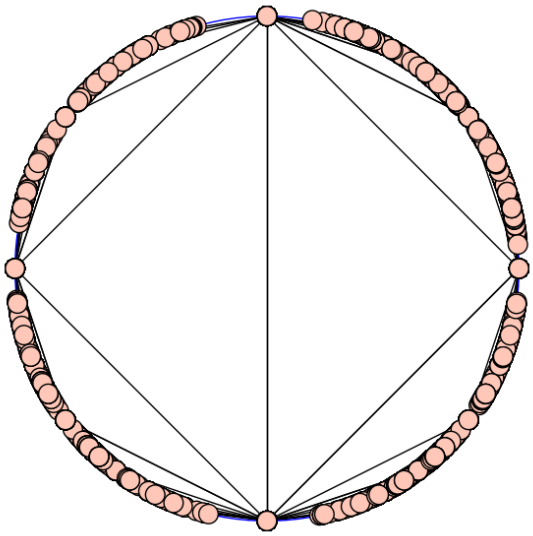

FIGURE 5.4.4. (Example 1) Projection of P_2 onto the Klein model

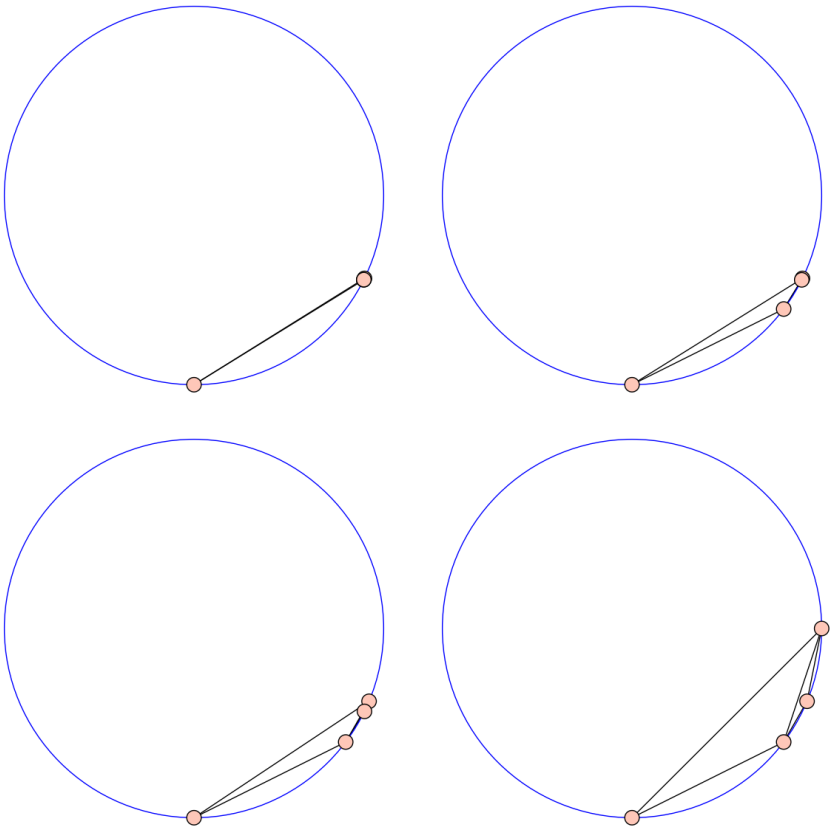

FIGURE 5.4.5. (Example 2) Projection of polyhedra P_i/Γ onto the Klein model. The order in which the polyhedra were visited by the edge flipping algorithm is top left, top right, bottom left, bottom right.

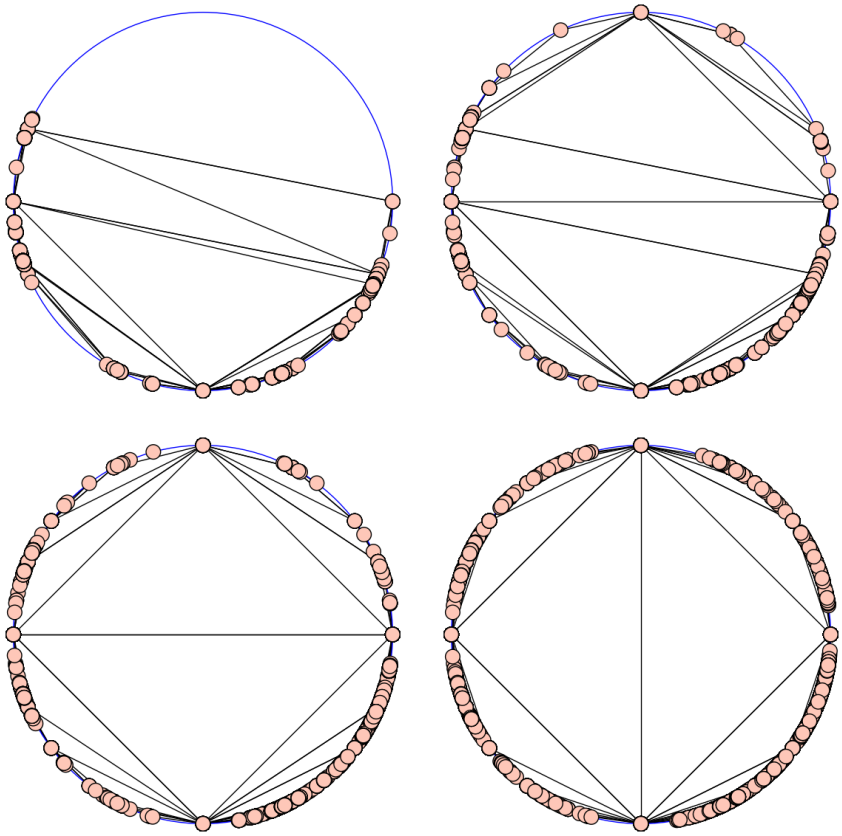

FIGURE 5.4.6. (Example 2) Projection of polyhedra P_i onto the Klein mode. The order in which the polyhedra were visited by the edge flipping algorithm is top left, top right, bottom left, bottom right.

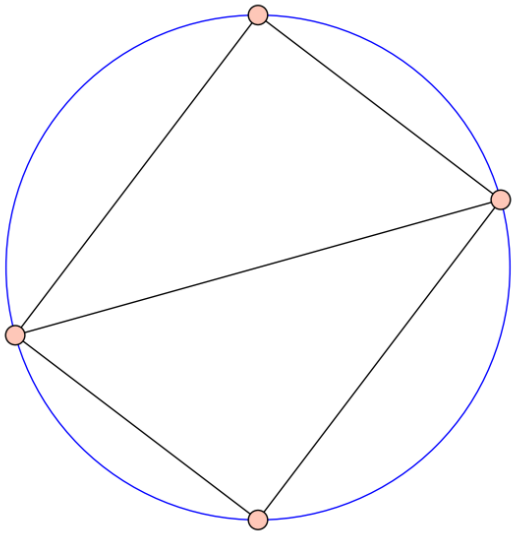

FIGURE 5.4.7. (Example 3) A hyperbolic structure close to the critical point in Figure 5.4.8.

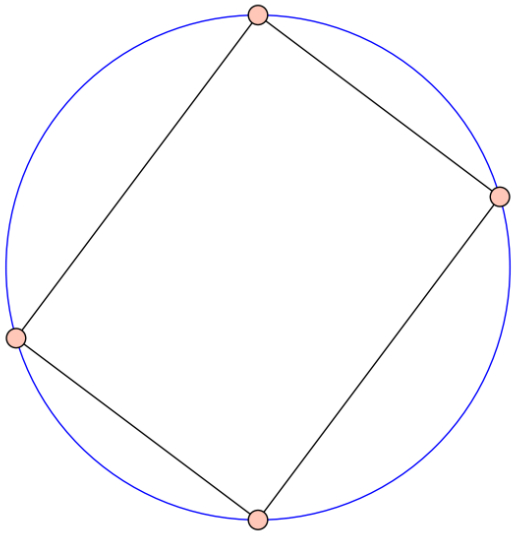

FIGURE 5.4.8. (Example 3) A critical point of the parameter space, where the canonical cell decomposition has one cell.

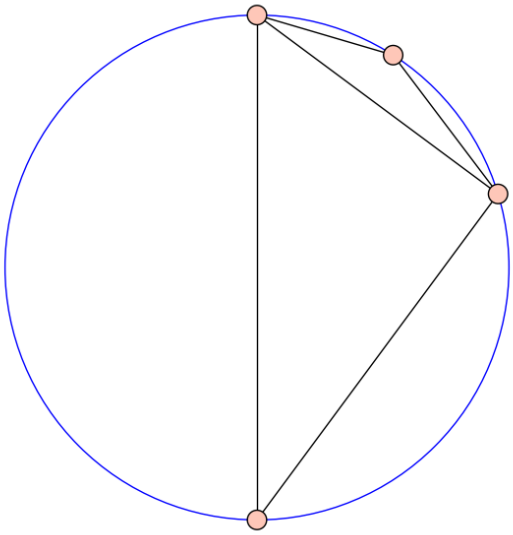

FIGURE 5.4.9. (Example 3) A hyperbolic structure close to the critical point in Figure 5.4.8.

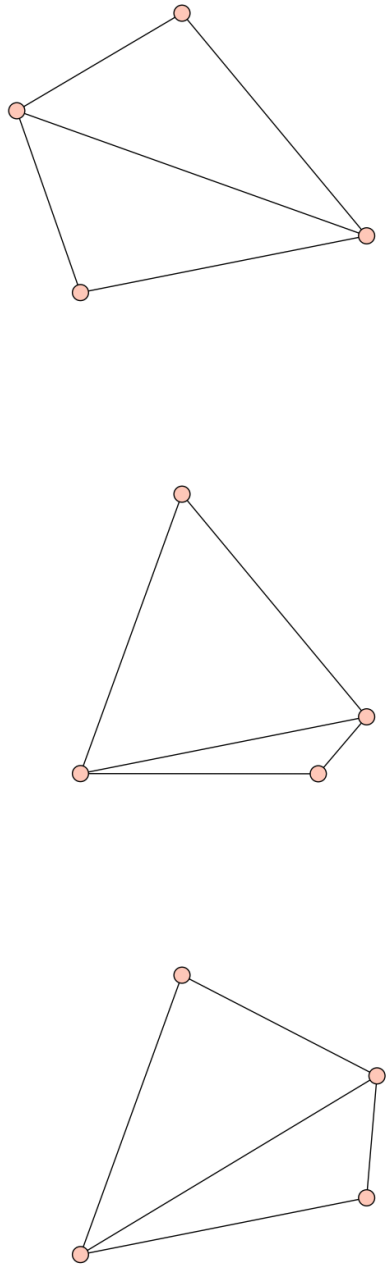

FIGURE 5.4.10. (Example 4) Projections of Γ-invariant polyhedra visited by Algorithm 1 for the projective structure (8,2,3,7,1,2).

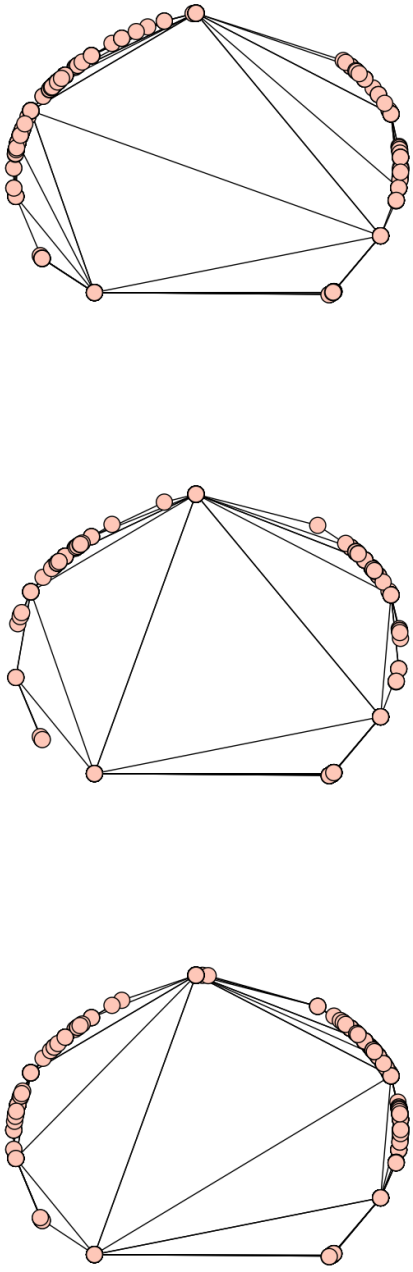

FIGURE 5.4.11. (Example 4) Projections of Γ-invariant polyhedra visited by Algorithm 1 for the projective structure (8,2,3,7,1,2xa points on $\partial\Omega$shown

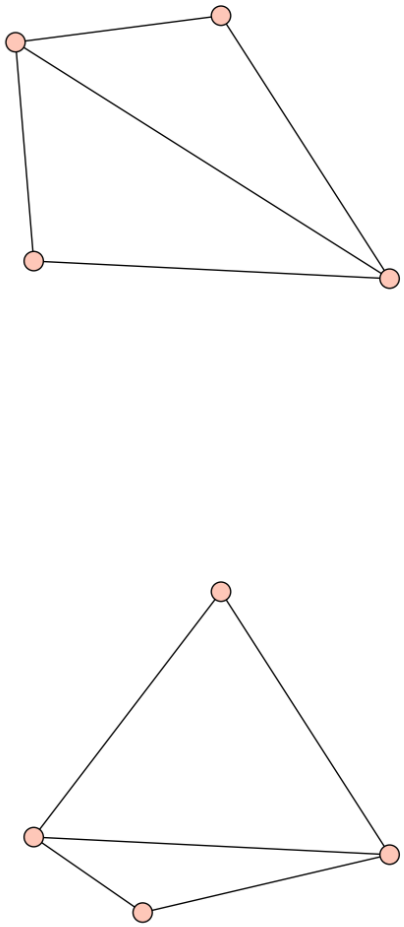

FIGURE 5.4.12. (Example 5) The intermediate cell decompositions for the projective structure corresponding to $(4,7,\frac{3}{2},1,1,7)$

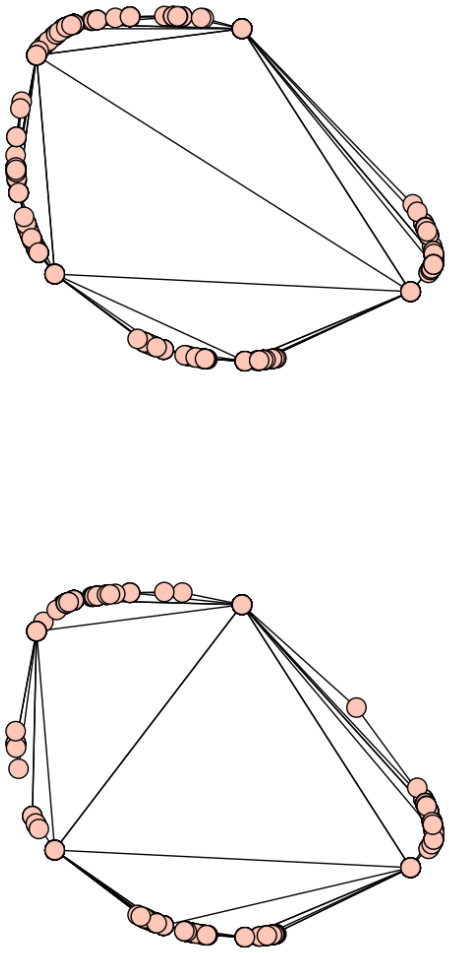

FIGURE 5.4.13. (Example 5) The intermediate cell decompositions for the projective structure corresponding to $(4,7,\frac{3}{2},1,1,7)$

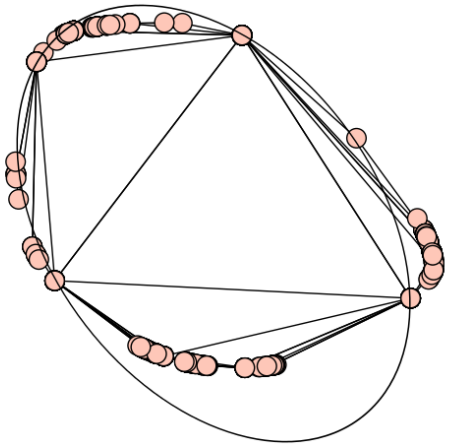

FIGURE 5.4.14. (Example 5) There is no ellipse passing through all points on the boundary ∂D.

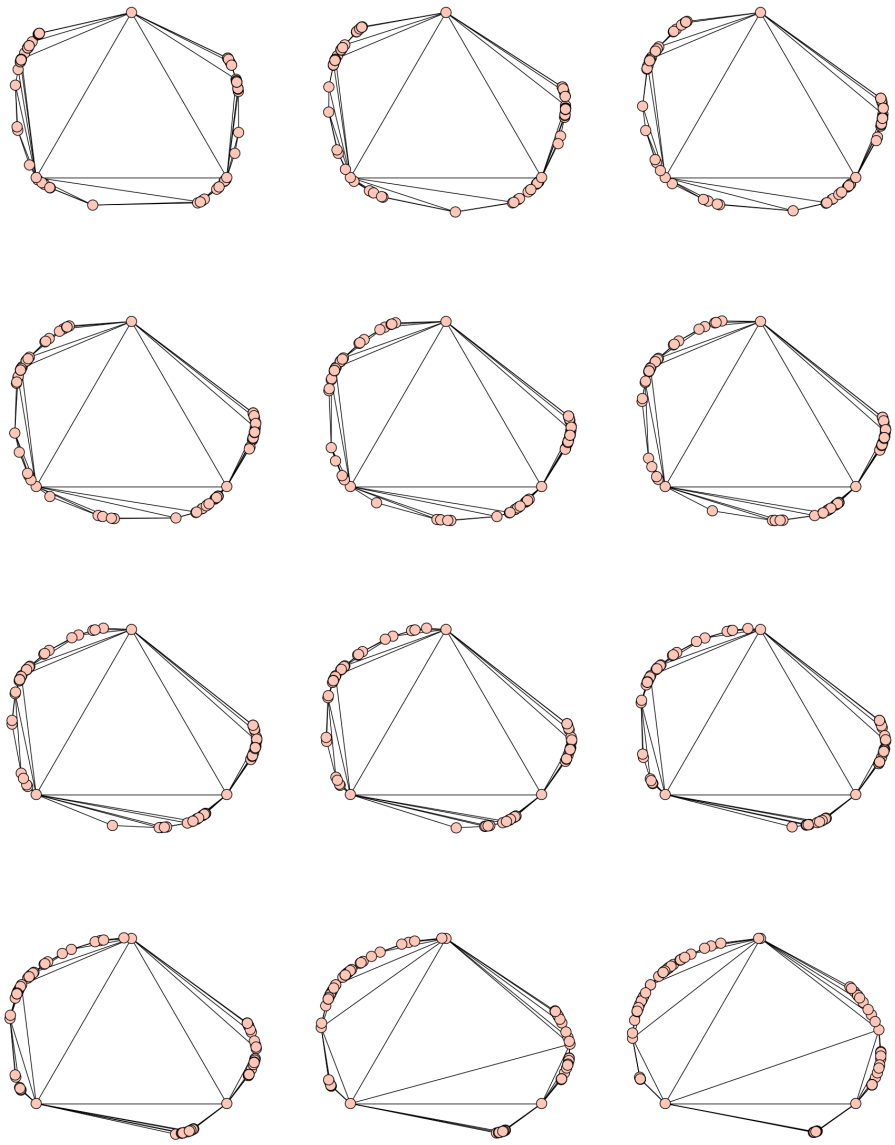

FIGURE 5.4.15. (Example 6) The cell decompositions corresponding to 12 evenly spaced points along the segment joining $(8, 2, 7, 2, 4, 2), (8, 2, 3, 7, 1, 2) \in Q$.

References

[BE88] Brian H Bowditch and David BA Epstein, *Natural triangulations associated to a surface*, Topology **27** (1988), no. 1, 91–117.

[Bon09] Francis Bonahon, *Low-dimensional geometry*, Student Mathematical Library, vol. 49, American Mathematical Society, Providence, RI; Institute for Advanced Study (IAS), Princeton, NJ, 2009, From Euclidean surfaces to hyperbolic knots, IAS/Park City Mathematical Subseries. MR 2522946 (2011f:57024)

[Bur14] Benjamin A Burton, *The cusped hyperbolic census is complete*, arXiv preprint arXiv:1405.2695 (2014).

[Bus10] Peter Buser, *Geometry and spectra of compact Riemann surfaces*, Modern Birkhäuser Classics, Birkhäuser Boston, Inc., Boston, MA, 2010, Reprint of the 1992 edition. MR 2742784 (2011i:58047)

[CGHN00] David Coulson, Oliver A Goodman, Craig D Hodgson, and Walter D Neumann, *Computing arithmetic invariants of 3-manifolds*, Experimental Mathematics **9** (2000), no. 1, 127–152.

[CHW99] Patrick Callahan, Martin Hildebrand, and Jeffrey Weeks, *A census of cusped hyperbolic 3-manifolds*, Mathematics of Computation of the American Mathematical Society **68** (1999), no. 225, 321–332.

[CL13] Daryl Cooper and Darren Long, *A generalization of the epstein-penner construction to projective manifolds*, arXiv preprint arXiv:1307.5016 (2013).

[CLT11] Daryl Cooper, Darren Long, and Stephan Tillmann, *On convex projective manifolds and cusps*, arXiv preprint arXiv:1109.0585 (2011).

[EP+88] David BA Epstein, Robert C Penner, et al., *Euclidean decompositions of noncompact hyperbolic manifolds*, Journal of Differential Geometry **27** (1988), no. 1, 67–80.

[G+90] William Goldman et al., *Convex real projective structures on compact surfaces*, J. Differential Geom **31** (1990), no. 3, 791–845.

[GCTH13] Mingcen Gao, Thanh-Tung Cao, Tiow-Seng Tan, and Zhiyong Huang, *Flip-flop: convex hull construction via star-shaped polyhedron in 3d*, Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, ACM, 2013, pp. 45–54.

[Har86] John L Harer, *The virtual cohomological dimension of the mapping class group of an orientable surface*, Inventiones mathematicae **84** (1986), no. 1, 157–176.

[HRS12] Craig D Hodgson, J Hyam Rubinstein, and Henry Segerman, *Triangulations of hyperbolic 3-manifolds admitting strict angle structures*, Journal of Topology (2012), jts022.

[KM13] Alexander Kolpakov and Bruno Martelli, *Hyperbolic four-manifolds with one cusp*, Geometric and Functional Analysis **23** (2013), no. 6, 1903–1933.

[Lac00] Marc Lackenby, *Taut ideal triangulations of 3-manifolds*, arXiv preprint math/0003132 (2000).

[Law77] Charles L Lawson, *Software for c1 surface interpolation. mathematical software iii (john˜ r. rice, editor), pages 161–194*, 1977.

[Mar10] Ludovic Marquis, *Espace des modules marqués des surfaces projectives convexes de volume fini*, Geom. Topol **14** (2010), no. 4, 2103–2149.

[Pen87] Robert C Penner, *The decorated teichmüller space of punctured surfaces*, Communications in Mathematical Physics **113** (1987), no. 2, 299–339.

[Rat06] John G. Ratcliffe, *Foundations of hyperbolic manifolds*, second ed., Graduate Texts in Mathematics, vol. 149, Springer, New York, 2006. MR 2249478 (2007d:57029)

[Ser99] Caroline Series, *Lectures on pleating coordinates for once punctured tori*, Hyperbolic Spaces and Related Topics, RIMS Kokyuroku **1104** (1999), 30–108.

[SW95] Makoto Sakuma and Jeffrey R Weeks, *The generalized tilt formula*, Geometriae Dedicata **55** (1995), no. 2, 115–123.

[Thu97] William P. Thurston, *Three-dimensional geometry and topology. Vol. 1*, Princeton Mathematical Series, vol. 35, Princeton University Press, Princeton, NJ, 1997, Edited by Silvio Levy. MR 1435975 (97m:57016)

[Ush99] Akira Ushijima, *A canonical cellular decomposition of the teichmüller space of compact surfaces with boundary*, Communications in mathematical physics **201** (1999), no. 2, 305–326.

[Ush02] ———, *The tilt formula for generalized simplices in hyperbolic space*, Discrete and Computational Geometry **28** (2002), no. 1, 19–27.

[Wee93] Jeffrey R Weeks, *Convex hulls and isometries of cusped hyperbolic 3-manifolds*, Topology and its Applications **52** (1993), no. 2, 127–149.

[Yos01] Han Yoshida, *Epstein–penner decompositions and thrice punctured spheres in hyperbolic 3-manifolds*, Geometriae Dedicata **87** (2001), no. 1-3, 1–16.