PL Embeddings of Circle Quotients and Continued Fractions

Richard Evan Schwartz *

November 16, 2002

1 Introduction

The string tensions in tennis rackets usually reside within a fairly narrow range, though some players string their rackets far above this range. For the most part this practice is prompted by the player's mistaken idea that an unusually tight string-job will enhance his or her performance. In this paper we address the question: What does a tennis racket look like if it is strung so tightly that the individual strings collapse into points?

This inquiry is unlikely to have a practical application, so we will save money by considering not the abominations of wood and graphite produced by an actual experiment but rather the theoretical objects produced by the following construction:

1. Let Γ be a finite union of closed geodesics on a hyperbolic surface Σ.
2. Let $\widetilde{\Gamma}$ be the lift of Γ to the universal cover (hyperbolic plane) $\boldsymbol{H}^2 = \widetilde{\Sigma}$. Let S^1 be the ideal boundary of $\boldsymbol{H}^2$.
3. Let $Q(\Sigma, \Gamma)$ be the quotient of S^1 by the relation that two points are equivalent if they are the endpoints of a geodesic $\widetilde{\gamma} \in \widetilde{\Gamma}$.

The space $Q(\Sigma, \Gamma)$, which we call a *circle quotient*, is an idealized version of a tennis racket strung at infinite tension.

* Supported by N.S.F. Research Grant DMS-0072607

The case when Γ is the union of one or two simple closed geodesics arises in connection with quasi-Fuchsian groups. Suppose ρ_n is a sequence of quasi-Fuchsian representations of the fundamental group $\pi_1(\Sigma)$, such that the words associated to the geodesics in Γ become parabolic as $n \to \infty$. Then the limit set of the limit group $\rho_\infty(\pi_1(\Sigma))$ is typically $Q(\Sigma, \Gamma)$. See [**T**] for a general discussion about the degeneration of quasi-Fuchsian groups.

In [**S1**] we studied some examples of $Q(\Sigma, \Gamma)$, when Σ is cusped (i.e. non-compact, finite volume) and Γ is a closed geodesic with some self-intersections. The motivating example arises in connection with complex hyperbolic deformations of the ideal triangle group [**S2**]. For each of the examples in [**S1**] (which are typically not planar) we give an embedding $Q \hookrightarrow \boldsymbol{R}^3$ which intertwines the action of $\pi_1(\Sigma)$ with a group of PL maps of the embedded set. (A PL map of a set is a self-homeomorphism of that set which extends to be piecewise linear in an open neighborhood of it.) We call such an embedding a *PL embedding*.

In this paper we give PL embedding in the general case:

Theorem 1.1 *Suppose that Σ is cusped and Γ is a finite union of closed geodesics on Σ. Then there is a PL embedding $Q(\Sigma, \Gamma) \hookrightarrow \boldsymbol{R}^d$, for some d.*

The dimension d increases unboundedly with the ratio length(Γ)/area(Σ). In general the embeddings we get are of such high dimension that they are impossible to picture. However, we think that in some nontrivial cases the embedding is sufficiently concrete that it really gives a good picture.

One case of interest in our construction occurs when $\widetilde{\Gamma}$ is invariant under the action of the modular group $PSL_2(\boldsymbol{Z})$ (and we set $S^1 = \boldsymbol{R} \cup \infty$, the ideal boundary of the hyperbolic upper half plane.) The construction is the cleanest when $\widetilde{\Gamma}$ is the lift of a smooth closed geodesic on the modular orbifold $\boldsymbol{H}^2/PSL_2(\boldsymbol{Z})$. We call this the *minimal case*. We now describe the 4-step construction of the PL embedding in the minimal case.

Step 1: Let ϕ be an endpoint of some geodesic of $\widetilde{\Gamma}$. Let $...n_0, n_1, n_2, ...$ be the periodic part of the continued fraction expansion of ϕ. We convert this infinite periodic sequence to a binary string with $...n_0$ 0's then n_1 1's then n_2 0's, ... (This is the *cutting sequence* in the sense of [**S**, Th. 5.18].) Let $\{b_k\}_{k=1}^s$ be the minimal finite portion of this binary string which generates the whole string. For instance, given the periodic continued fraction ...131131131... leads to the finite sequence 0111010001. In this example $s = 10$.

Step 2: Let $V_s = \{0_0, 1_0, 2_0, ..., s_0, 0_1, 1_1, 2_1, ..., s_1\}$. We define an involution $I: V_s \to V_s$ by the following rules: $I(0_0) = 0_1$ and

$$I(k_{b_k}) = (k+1)_{1-b_{k+1}} \qquad s = 1, ..., k. \tag{1}$$

When $k = s$ we interpret $k+1$ as 1. For instance, if $\{a_k\} = 110101$ then I has the action $(0_0, 1_1, 2_1, 3_0, 4_1, 5_0, 6_1) \to (0_1, 2_0, 3_1, 4_0, 5_1, 6_0, 1_0)$.

Step 3: We construct a triple T_0 of d-simplices $\Delta_1, \Delta_2, \Delta_3$, with pairwise disjoint interiors and empty triple intersection, such that $\partial\Delta_i \cap \partial\Delta_j$ is the convex hull of $s+1$ common vertices, for all $i \neq j$. We label the vertices of Δ_1, Δ_2, and Δ_3 by the elements of V_s such that each vertex gets a different label, and the m_0-vertex of Δ_j coincides with the m_1-vertex of Δ_{j+1} for $m = 0, 1, ..., s$. Here j is taken mod 3. Likewise we construct a triple T' of d-simplices $\Delta_1', \Delta_2', \Delta_3'$ such that $\Delta_1', \Delta_2' \subset \Delta_3'$ and the boundaries have the same intersections as the boundaries of the simplices in T. We label the vertices of Δ_j' just as we labelled the vertices of Δ_j. We set $\Delta' = \Delta_3'$.

Step 4: Suppose T_m has been constructed, and consists of a finite union of labelled simplices. For each Δ of T_m we choose the affine map $\phi: \Delta' \to \Delta$ such that ϕ maps the i_j-vertex of Δ' to the $I(i_j)$-vertex of Δ, for all $i_j \in V_s$. We define $T_{m+1} = \bigcup \phi(\Delta_1' \cup \Delta_2')$. The union is taken over all simplices Δ of T_m. We let $X(I) = \bigcap T_m$. A *rational vertex* of X is a vertex of some simplex of $\bigcup T_m$ which is labelled by a 0_0 or 0_1. An *irrational vertex* of X is some other vertex of a simplex of $\bigcup T_m$.

Theorem 1.2 *There is a PL embedding $Q(\widetilde{\Gamma}) \to X$, induced by a continuous surjection $R \cup \infty \to X$ which maps the rational points (i.e. the modular group cusps) bijectively onto the rational vertices, and the endpoints of geodesics in $\widetilde{\Gamma}$, in two-to-one fashion, onto the irrational vertices.*

There is a similar result when $\widetilde{\Gamma}$ is the lift of a closed geodesic which contains the order 2 orbifold point on the modular orbifold. In this case, Steps 1 and 2 are slightly modified, as we will explain in the remark at the end of §5. The circle quotient which arises in connection with complex hyperbolic ideal triangle groups corresponds to the finite sequence $\{1\}$ in the modified construction, and is discussed in the same remark. In this case, our embedding is implemented by spherical CR geometry—i.e. the geometry of S^3 which is adapted to the automorphism group $PU(2,1)$.

It seems that we could also analyze the case when our pattern of geodesics is based on a closed surface, though we do not yet know how to do this.

All the circle quotients we consider in this paper have embeddings into $\boldsymbol{R}^3$, though in general the embeddings seem quite "distorted". In a future paper we hope to quantify this distortion, and takes steps towards answering the question: Given a Lie group G, acting on a homogeneous space X, which circle quotients arise as limit sets of discrete subgroups of G acting on X? An effective answer to this question would provide a new tool for studying deformations of surface groups into Lie groups.

This paper is organized as follows. In §2 we prove Theorem 1.1 modulo an auxilliary result, the Correspondence Lemma. In §3 we prove the Correspondence Lemma. In §4 we analyze modular circle quotients, and prove a more abstract version of Theorem 1.2. In §5 we investigate the connection between the abstract construction in §4 and continued fractions, thereby proving Theorem 1.2 as stated.

I would like to thank the I.H.E.S. for their great hospitality during the writing of this paper.

2 General Construction

2.1 The Set Up

In this chapter we will assume that Σ is a cusped hyperbolic surface. This means that the universal covering group Λ acts freely on $\boldsymbol{H}^2 = \widetilde{\Sigma}$. Our construction depends on some auxilliary choices for Λ. Let T be a Λ-invariant tiling of $\boldsymbol{H}^2$ by ideal triangles. To find T we first pick a finite sided fundamental polygon for Λ and, then triangulate this polygon, then extend the triangulation equivariantly. Let E denote the set of edges of the tiling T.

Let T^* denote the tree dual to T. We choose some triangle t_0 of T and call the corresponding node of T^* the *initial node*. We direct T^* outwards from its initial node. An edge e^* of T^* has a *tail vertex* and a *head vertex*, the names being chosen so that e^* is directed from its tail to its head. Given two edges $e_1, e_2 \in E$ we write $e_1 \to e_2$ iff e_2 is not an edge of t_0 and the head vertex of e_1^* coincides with the tail of e_2^*. Here e_j^* is the edge of T^* which corresponds to e_j. We write $e \Rightarrow e'$ iff there is a finite chain $e \to e_1 \to \ldots \to e_n \to e'$.

Let γ be a geodesic of $\widetilde{\Gamma}$. It is well known that γ cannot be asymptotic to an edge of $e \in E$. Hence γ either crosses e or remains a uniformly bounded distance from e. We let $\widetilde{\Gamma}(e)$ denote the set of geodesics in $\widetilde{\Gamma}$ which cross e.

Lemma 2.1 *There are only finitely many geodesics in $\widetilde{\Gamma}(e)$.*

Proof: There exist horodisk neighborhoods of the cusps on Σ, which are disjoint from the geodesics in Γ. Lifting these neighborhoods to $\boldsymbol{H}^2$, we get a union H of horodisks, one per vertex of the tiling $\Lambda(F)$, which is disjoint from $\widetilde{\Gamma}$. Any $\gamma \in \widetilde{\Gamma}(e)$ must cross e at a point $e \cap \gamma \in e - H$. There is a uniform lower bound to the distance between any two points $e \cap \gamma_1$ and $e \cap \gamma_2$. The set $e - H$ is bounded and hence there are finitely many points of the form $e \cap \gamma$. To finish the proof, note that there are only finitely many geodesics of $\widetilde{\Gamma}$ containing any point of $\boldsymbol{H}^2$. ♠

The cardinality $|\widetilde{\Gamma}(e)|$ only takes on finitely many values. Define

$$D = \max_e |\widetilde{\Gamma}(e)| + 2; \qquad d = D - 1; \qquad n(e) = D - 2 - |\widetilde{\Gamma}(e)|. \qquad (2)$$

If $n(e) > 0$ we let $p_1, \ldots, p_n$ be $n = n(e)$ points on e. These points (whose specific locations play no role in our construction) can be chosen arbitrarily, subject to the constraint that they do not lie on any other $e' \in E$. We call these points the *padding points*.

2.2 The Correspondence

For each edge $e \in E$ we assign a d-dimensional simplex $\Delta(e) \subset \boldsymbol{R}^d$. Our main idea is to label the vertices of $\Delta(e)$ *by geometric objects*. This means that we associate to each vertex a certain geometric object, which we think of as its label. Here is our labelling convention:

1. Label the first vertex of $\Delta(e)$ by the tail endpoint of e.
2. Label the next $|\widetilde{\Gamma}(e)|$ vertices by the geodesics in $\widetilde{\Gamma}(e)$.
3. Label the last vertex by the head endpoint of e.
4. Label the remaining $n(e)$ vertices by the padding points of e.

We say that a *correspondence* is the assignment of simplices in $\boldsymbol{R}^d$ to edges of E, subject to the labelling scheme just given.

Let e_1 and e_2 be two edges of E. Suppose $e_2 = \gamma(e_1)$ for some $\gamma \in \Lambda$ and suppose also that there is an affine map T such that $T(\Delta(e_1)) = \Delta(e_2)$. We say that T and γ are *intertwined* if the following condition holds. If a vertex v of $\Delta(e_1)$ is labelled by the geometric object K then the vertex $T(v)$ of $\Delta(e_2)$ is labelled by the geometric object $\gamma(K)$.

Lemma 2.2 (Correspondence Lemma) *The correspondence can be chosen so that*

1. *The interior of $\Delta(e')$ is contained in $\Delta(e)$ iff $e \Rightarrow e'$. Otherwise $\Delta(e)$ and $\Delta(e')$ have disjoint interiors.*
2. *$\partial\Delta(e) \cap \partial\Delta(e')$ is the convex hull of the set of vertices labelled by the geometric objects associated to both $\Delta(e)$ and $\Delta(e')$.*
3. *Suppose $e_1 \Rightarrow e_2$ and $f_1 \Rightarrow f_2$. If there is a $\gamma \in \Lambda$ such that $\gamma(e_1) = f_1$ and $\gamma(e_2) = f_2$ then there is an affine map T, intertwined with γ, such that $T(\Delta(e_1)) = \Delta(f_1)$ and $T(\Delta(e_2)) = \Delta(f_2)$.*

Moreover, for any $\epsilon > 0$ there are only finitely many simplices of the form $\Delta(e)$ having diameter greater than ϵ.

2.3 The Limit Object

Let Δ_∞ denote the set of points in $\boldsymbol{R}^d$ which are contained in infinitely many different simplices of the form $\Delta(e)$. We will construct a homeomorphism from $Q(\Sigma,\Gamma)$ to Δ_∞.

There are two kinds of points in S^1. First, there is a countable list of tiling vertices. Second, there are the remaining *generic points*. We consider the tiling vertex points first. If x is a vertex of a tiling, we pick any edge $e \in E$ which contains x as an endpoint. We define $\phi(x)$ to be the corresponding vertex of $\Delta(e)$. It follows immediately from the second condition of the Correspondence Lemma that the definition of $\phi(x)$ is independent of the choice of e.

Each edge e of T divides $\boldsymbol{H}^2$ into two halfspaces. The dual edge e^* points into one of these halfspaces, which we call *the right halfspace*.

If $x \in S^1$ is a generic point, there is an infinite sequence of edges

$$e_1 \to e_2 \to e_3...$$

converging to x. Moreover, e is in the right halfspace of e_j for all j. We define

$$\phi(x) = \bigcap \Delta(e_n) \tag{3}$$

Using the correspondence lemma, and the fact that T^* is a tree, the simplices on the right hand side of Equation 3 are nested and their intersection is a single point.

Lemma 2.3 *ϕ is a homeomorphism.*

Δ_∞ is a closed and bounded subset of $\boldsymbol{R}^d$, and hence is a compact Hausdorff space. The quotient $Q(\Sigma,\Gamma)$ is a compact space. It is well known that a continuous bijection from a compact space to a Hausdorff space is a homeomorphism. Hence, it suffices to show that ϕ is a continuous bijection.

Given any ϵ there is some N such that any nested intersection of the form $\Delta_1 \cap ... \cap \Delta_n$ has diameter at most ϵ. This follows immediately from the Correspondence Lemma. Given N there is some δ with the following property. If $x, y \in S_1$ have the property that $d(x,y) < \delta$ then there is a string of N edges $e_1 \to ... \to e_n$ such that x and y are both in the right halfspace of $\boldsymbol{H}^2 - e_j$, for $j = 1,...,N$. Thus means that $\phi(x)$ and $\phi(y)$ are both contained a nested intersection of the type just considered, where

$\Delta_j = \Delta(e_j)$, and hence are within ϵ of each other. This proves that ϕ is continuous.

Suppose that x_1 and x_2 are two distinct points on S^1. Then there are edges e_1 and e_2 such that the right half spaces H_1 and H_2, respectively bounded by e_1 and e_2 have disjoint closures, and respectively contain x_1 and x_2. From the conditions of the Correspondence Lemma this means that $\Delta(e_1)$ and $\Delta(e_2)$ are disjoint. Since $\phi(x_j) \in \Delta(e_j)$ we see that $\phi(x_1) \neq \phi(x_2)$. This proves that ϕ is injective.

To prove that ϕ is surjective, let $y \in \Delta_\infty$ be a point. There is an infinite sequence of simplices $\{\Delta(e_j)\}$, corresponding to a sequence $e_1 \to e_2 \to e_3 ...$ such that $y \in \Delta(e_j)$ for all j. The sequence $\{e_j\}$ has a unique accumulation point $x \in S^1$. By construction $\phi(x) = y$. This proves that ϕ is surjective and hence bijective. ♠

We say that a *PL automorphism* of a set $K \subset \boldsymbol{R}^d$ is a self-homeomorphism of K which extends to be piecewise-linear in a neighborhood of K.

Lemma 2.4 *There is a group Ω of PL homeomorphisms of Δ_∞ which is ϕ-conjugate to the action of Λ on $Q(\Sigma, \Gamma)$.*

Proof: Recall that E is the set of edges of the Λ-invariant tiling T. Let $R(e) \subset S^1$ denote the closed boundary of the right halfspace of e. Say that a *distinguished arc* is an arc of the form $R(e)$ for some edge $e \in E$. Let $\gamma \in \Lambda$ be some element. Since γ respects the tiling T, we can write $S^1 = R_1 \cup ... \cup R_n$, where

1. The distinguished arcs R_j have pairwise disjoint interiors.
2. $R'_j = \gamma(R_j)$ is another distinguished arc (probably not on the original list.)

Let e_j be the edge in E corresponding to R_j. Let $e'_j = \gamma(e_j)$. Let $E(\gamma)$ denote all those edges e such that $R_j \Rightarrow e$ for some j. Let $E'(\gamma)$ denote all those edges e such that $R'_j \Rightarrow e$.

Let $\Delta_j = \Delta(e_j)$ and also let $\Delta'_j = \Delta(e'_j)$. Given any j there is a canonical affine map from Δ_j onto Δ'_j. To define this map it suffices to consider the vertices. If x is a vertex of Δ_j then there is some geometric object g_x which labels x. We map x to the point x', which is labelled by the geometric object

$\gamma(g_x)$. The two maps $\Delta_j \to \Delta'_j$ and $\Delta_k \to \Delta'_k$ are consistent on their common domain of definition.

We have now defined a PL map $\omega : \bigcup \Delta_j \to \bigcup \Delta'_j$. If follows immediately from Condition 3 of the Correspondence Lemma that

$$\omega(\Delta(e)) = \Delta(\gamma(e)) \qquad \forall e \in E(\gamma). \tag{4}$$

Equation 4 easily implies that the restriction of ω to Δ_∞ is ϕ-conjugate to the action of γ on S^1.

The set $\bigcup \Delta_j$ contains Δ_∞ but does not contain an open neighborhood of Δ_∞. However, given that $\bigcup \Delta_j$ is a finite union of simplices, it is easy to extend ω to be piecewise linear in a neighborhood of Δ_∞. In short, ω is a PL automorphism of Δ_∞ which is ϕ-conjugate to γ. ♠

3 The Correspondence Lemma

3.1 Triples of Simplices

In this section we establish two technical results. Let $H(S)$ denote the convex hull of a set $S \subset \boldsymbol{R}^d$.

Lemma 3.1 *Suppose $s_1, s_2, s_3 \in \boldsymbol{N}$ satisfy $s_i + s_{i+1} \leq d+1$ for $i = 1, 2, 3$. Then there exist d-dimensional simplices $\Delta_1, \Delta_2, \Delta_3$ such that*

1. *The interiors of the Δs are pairwise disjoint.*
2. *$\partial\Delta_{j-1} \cap \partial\Delta_{j+1}$ is the convex hull of exactly s_j common vertices.*
3. *$\Delta_1 \cap \Delta_2 \cap \Delta_3 = \emptyset$.*

Here indices are taken mod 3.

Proof: Let $d_j = s_{j-1} + s_{j+1}$. Suppose that we could construct, for $j = 1, 2, 3$, a d_j-dimensional simplex Δ'_j, such that Δ'_1, Δ'_2, and Δ'_3 satisfy the conclusions of the lemma. We could then construct Δ_j by adding $d - d_j$ new vertices to Δ'_j. If we choose the new vertices in general position, close to $\partial\Delta'_j$, and far from the vertices of Δ'_j, then we can guarantee that $\Delta_{j-1} \cap \Delta_{j+1} = \Delta'_{j-1} \cap \Delta'_{j+1}$. Hence, it suffices to construct the Δ'_j.

Let $\sigma_1, \sigma_2, \sigma_3$ be the vertices of an equilateral triangle in $\boldsymbol{R}^d$. Let Π_j be the hyperplane containing σ_j and perpendicular to $\overline{\sigma_{j-1}\sigma_{j+1}}$. We choose some small $\epsilon > 0$. Let S_j be a regular simplex which is the union of s_j points contained in Π_j and exactly ϵ from σ_j. Define

$$\Delta'_j = H(S_{j-1} \cup S_{j+1}).$$

Below we will show that Δ'_j is a d_j-simplex. Assume this for the moment.

$\boldsymbol{R}^d - \Pi_j$ consists of two half-spaces. One of these half-spaces contains σ_{j-1} and the other half-space contains σ_{j+1}. Once ϵ is sufficiently small, all the vertices of $\Delta'_{j-1} - S_j$ are contained in the first half space and all the vertices of $\Delta'_{j+1} - S_j$ are contained in the second half space. Hence $\Delta_{j-1} \cap \Delta_{j+1} = H(S_j)$. For ϵ small we have $H(S_i) \cap H(S_j) = \emptyset$ and hence $\Delta'_1 \cap \Delta'_2 \cap \Delta'_3 = \emptyset$.

To finish our proof we need to show that Δ'_j is a d_j-simplex once ϵ is sufficiently small. Δ'_j is the convex hull of $d_j + 1$ points. It suffices to prove that none of these points is a convex combination of the others. Suppose some $x \in S_{j-1}$ is a convex combination of the remaining points. Since x is an extreme point of $H(S_{j-1})$ it must happen that x is contained on a line L joining a point in $H(S_{j-1})$ to a point in $H(S_{j+1})$. Since L contains two distinct points of Π_{j-1}, this entire line must lie in Π_{j-1}. Once ϵ is sufficiently small, L is too nearly parallel to the line $\overline{\sigma_{j-1}\sigma_{j+1}}$ for it to lie in Π_{j-1}. This contradiction shows that every point of S_{j-1} is an extreme point of Δ'_j. The same argument works for points in S_{j+1}. ♠

Lemma 3.2 *Suppose $s_1, s_2, s_3 \in \boldsymbol{N}$ satisfy $s_i + s_{i+1} \leq d+1$ for $i = 1, 2, 3$. Then there exist d-dimensional simplices $\Delta_1, \Delta_2, \Delta_3$ such that*

1. *$\Delta_1, \Delta_2 \subset \Delta_3$.*
2. *$\partial\Delta_{j-1} \cap \partial\Delta_{j+1}$ is the convex hull of exactly s_j common vertices.*
3. *$\partial\Delta_1 \cap \partial\Delta_2 \cap \partial\Delta_3 = \emptyset$.*

Here indices are taken mod 3.

Proof: Let Π_1 and Π_2 be parallel hyperplanes. Let S_j be s_j points in Π_j which comprise the vertices of a regular simplex. Choose s'_1 so that $s'_1 + s_2 = d+1$. Let $S'_1 \supset S_1$ be s'_1 points in Π_j which comprise the vertices of a regular simplex.

Let $\Delta_3 = H(S'_1 \cup S_2)$. It is easy to see that Δ_3 is a d-dimensional simplex. Let Π_3 be the hyperplane parallel to Π_1 and Π_2 and halfway between them. Note that $\Pi_3 \cap \Delta_3$ contains an open set. Let S_3 be s_3 points, comprising the vertices of a regular simplex, and all contained in the interior of $\Pi_3 \cap \Delta_3$. Let $\Delta'_1 = H(S_1 \cup S_3)$ and $\Delta'_2 = H(S_1 \cup S_3)$. An argument similar to the one we gave in the previous lemma shows that Δ'_1, Δ'_2 and Δ_3 satisfy the three listed conditions.

If $\dim(\Delta'_1) < d$ then we add the requisite number of vertices to Δ'_1, as we did in the previous lemma, to create the d-dimensional simplex Δ_1. We do the same for Δ_2. Once again, we add these new vertices in such a way as to retain the desired intersection properties. ♠

3.2 The Main Construction

Recall that t_0 is the distinguished triangle of our ideal triangulation T. Let e_1, e_2, e_3 be three edges bounding T_0. Let s_j be the number of geometric objects common to e_{j-1} and e_{j+1}, with indices taken mod 3. Note that there is no geometric object common to all three edges. This means that the triple (s_1, s_2, s_3) satisfies the hypotheses of Lemma 3.1. We construct three d-simplices according to Lemma 3.1 and then label them in the obvious way.

Each remaining triangle t of T has three edges one of which is distinguished. The distinguished edge of t is the one which is combinatorially closest to t_0. Just as we did for t_0, we can assign a triple of d-simplices to the edges of t. This time, however, we want to use Lemma 3.2 in place of Lemma 3.1. Furthermore, we want to make the assignment so that the simplex corresponding to the distinguished edge contains the other two simplices in its interior.

We can make our correspondence satisfy the equivariance part of the Correspondence Lemma by doing it inductively, and using the group action whenever possible. Suppose that we have make our assignment for all triangles which have a combinatorial distance at most n from t_0. Suppose t has combinatorial distance $n+1$ from t_0. Let e_1, e_2, e_3 be the edges of t, with e_1 being distinguished. By induction a simplex $\Delta(e_1)$ has already been defined. There are two cases to consider:

Case 1: Suppose there is some other edge e_1' such that $e_1' \Rightarrow e_1$ and $\gamma(e_1') = e_1$ for some $\gamma \in \Lambda$. Then e_1' is the distinguished edge of some triangle t', which is at most n combinatorial units from t_0. Let e_2' and e_3' be the other edges of t'. There is a unique affine map $T : \Delta(e_1') \to \Delta(e_1)$ which intertwines with γ. We define $\Delta(e_j) = T(\Delta(e_j'))$.

Case 2: There is no such edge e_1', as in Case 1. In this case we use Lemma 3.2 to construct the desired triple of simplices for t, with the only constraint being that $\Delta(e_1)$ coincides with the simplex already defined for e_1 at a previous stage of the construction. We then label the simplices in the obvious way.

It follows from the fact that T^* is a tree, and from the fact that the composition of affine maps is affine, that our construction is well defined and satisfies the first two conditions of the Correspondence Lemma.

3.3 Diameter Estimate

To finish our proof of the Correspondence Lemma we verify the diameter estimate for the correspondence we have constructed. Our proof is a generalization of the analogous proof in [**S1**].

If $e \Rightarrow e'$ we write $e \Rightarrow_n e'$ if the chain $e = e_1 \to \ldots \to e_n = e'$ (which is unique) has length exactly n.

Lemma 3.3 *There is some n with the following property: If $e \Rightarrow_n e'$ then $\partial\Delta(e) \cap \partial\Delta(e')$ is at most a single point.*

Proof: We can assume without loss of generality that e is one of the edges of F the fundamental polygon. There are two cases to consider. Suppose first that e and e' have a common endpoint. Then $\partial\Delta(e)$ and $\partial\Delta(e')$ share the common vertex labelled by this endpoint. As $n \to \infty$ the edge e' shrinks to the common endpoint of e and e'. But, there is a minimum distance $\delta > 0$ such that any of the finitely many geodesics of $\widetilde{\Gamma}(e)$ remain at least δ from either endpoint of e. Hence, for sufficiently large n, no geodesic of $\widetilde{\Gamma}$ which crosses e also crosses e'. This means that $\partial\Delta(e)$ and $\partial\Delta(e')$ share no other vertices.

In the second case, e and e' do not share a common vertex. As $n \to \infty$, the edge e' shrinks to a point on S^1. Two distinct geodesics of $\widetilde{\Gamma}$ to not share any common endpoints. Therefore, there is some $\delta > 0$ such that any two geodesics of $\widetilde{\Gamma}(e)$ are at least δ apart from each other on S^1. So, once n is sufficiently large, $\widetilde{\Gamma}(e)$ and $\widetilde{\Gamma}(e')$ share at most one common member. ♠

Let n be any number guaranteed by the previous lemma. Let $\mathcal{S}$ denote the set of pairs of simplices of the form $(\Delta(e), \Delta(e'))$, where $e \Rightarrow_n e'$. We say that two pairs (Δ_1, Δ_1') and (Δ_2, Δ_2') in $\mathcal{S}$ are *equivalent* if there is an affine transformation $T : \boldsymbol{R}^d \to \boldsymbol{R}^d$ such that $T(\Delta_1) = \Delta_2$ and $T(\Delta_1') = \Delta_2'$.

Lemma 3.4 *There are only finitely many equivalence classes in $\mathcal{S}$.*

Proof: Modulo the action of Λ there are only finitely many pairs of edges of the form (e, e') where $e \Rightarrow_n e'$. Our lemma follows immediately from this fact, and from the Condition 3 of the Correspondence Lemma. ♠

For each equivalence class in $\mathcal{S}$ we define a *model pair* (Δ, Δ'), affine equivalent to any member of the equivalence class, such that Δ is the unit regular simplex.

If the diameter estimate in the Correspondence Lemma is wrong, we can find a nested sequence $\Delta_1 \supset \Delta_2 \supset \Delta_3...$ such that $\bigcap \Delta_j$ is more than a single point. By taking an evenly spaced subsequence we can assume that (Δ_j, Δ_{j+1}) is a member of $\mathcal{S}$ for all j. At least one model pair (Δ, Δ') is represented for infinitely many different values of j.

Let λ'_j denote a line segment, contained in Δ_{j+1}, which has maximum length. Let λ_j be the larger line segment which is the intersection of Δ_j with the line containing λ'_j. Since $\bigcap \Delta_j$ does not shrink to a point, the length of λ'_j does not tend to 0. Moreover, since Δ_j and Δ_{j+1} converge to each other as $j \to \infty$, the lengths of λ_j and λ'_j converge to each other.

For an infinite collection of indices j there is an affine map T_j which takes the pair (Δ_j, Δ_{j+1}) to the model pair (Δ, Δ'). Let $\delta_j = T_j(\lambda_j)$ and $\delta'_j = T_j(\lambda'_j)$. An affine map respects ratios of distances on lines. Hence

$$\lim_{j \to \infty} \frac{\text{length}(\delta'_j)}{\text{length}(\delta_j)} = 1 \tag{5}$$

A maximum length segment in a simplex coincides with an edge of that simplex. Since Δ' has at most finitely many edges, we can assume that there is a single edge δ' of Δ' such that $\delta'_j = \delta'$ for all j (in our subsequence). This means that δ is independent of j as well. Equation 5 now says that $\delta' = \delta$. This means that $\partial\Delta$ and $\partial\Delta'$ share both endpoints of $\delta' = \delta$. This contradicts the definition of $\mathcal{S}$ and completes our proof.

4 Modular Circle Quotients

4.1 Basic Definitions

The modular group is $PSL_2(\boldsymbol{Z})$, acting on the upper half plane by linear fractional transformations. We set $\Lambda = PSL_2(\boldsymbol{Z})$. The *modular surface* is the quotient $\boldsymbol{H}^2/\Lambda$. This quotient is an orbifold, though it is finitely covered by a genuine surface. We say that a *modular pattern* is a Λ-invariant collection $\widetilde{\Gamma}$ of geodesics $\boldsymbol{H}^2$ which is the lift of a finite union of closed geodesics on some surface which finitely covers the modular surface.

Λ is the orientation preserving symmetry group of familiar tiling of the hyperbolic plane by ideal triangles, which we call *modular tiling*. The edges of the tiling are the geodesics whose endpoints have the form p_1/q_1 and p_2/q_2 where

$$\det \begin{bmatrix} p_1 & q_1 \\ p_2 & q_2 \end{bmatrix} = \pm 1 \tag{6}$$

Λ acts transitively on the set of ideal triangles and, moreover, transitively on the set $\overrightarrow{E}$ of oriented edges of the tiling.

There is a nice way to partition the set of ideal triangles in the modular tiling. Let F_0 be the ideal triangle with vertices $0 = 0/1$ and $1 = 1/1$ and $\infty = 1/0$. In general, suppose that F_m is a finite union of ideal triangles. We let $e_1, .., e_n$ be the geodesics of triangles in F_m which are not geodesics of triangles in F_{m-1}. We call these edges the *free edges*. (In general $n = 3 \times 2^m$.) Each e_j bounds an ideal triangle T_j' not in F_m. We let $F_{m+1} = \bigcup T_j'$. The union is taken over all the free edges of F_m. The sequence $\{F_n\}$ is the partition.

4.2 The Associated Involution

In the introduction we explained how to associate an involution to a minimal modular pattern. Here we give a general construction. In §5 we will see that the construction we give here specializes to the one we give in the introduction.

Let $\widetilde{\Gamma}$ be a modular pattern. Let D be as in Equation 2. That is, there are $D-2$ geodesics of $\widetilde{\Gamma}$ which cross each edge of the modular tiling. In this section we associate to $\widetilde{\Gamma}$ a certain involution defined on V_s, the D-element set defined in Step 2 of §1. In the next section we will use I to give a concrete embedding of the circle quotient $Q(\widetilde{\Gamma})$.

By symmetry $|\widetilde{\Gamma}(e)|$ is independent of $e \in E$. This means that there are no padding points, as defined in §2. Let s denote the number of geodesics common to two edges e_0 and e_1 which together bound an ideal triangle. By symmetry s is independent of the choice of e_0 and e_1. Since there are no padding points we have $D = 2s + 2$.

It is convenient to work with the oriented geodesics of $\overrightarrow{E}$. Let $\overrightarrow{e}$ be such a geodesic. There are unique oriented geodesics $\overrightarrow{e_0}$ and $\overrightarrow{e_1}$ which have the property that the triple $\overrightarrow{e}, \overrightarrow{e_0}, \overrightarrow{e_1}$ is the counterclockwise oriented boundary of an ideal triangle of the modular tiling. Since $D = 2(s+1)$, each geometric object associated to e is also associated to e_0 or e_1. For $i = 0, 1$ we let ${}_i(\overrightarrow{e})$ denote the geometric objects associated to both e and e_i.

For every edge $\overrightarrow{e}$ we order the objects of ${}_0(\overrightarrow{e})$ in some way, subject to the constraint that the head endpoint of $\overrightarrow{e}$, which necessarily belongs to ${}_0(\overrightarrow{e})$, is listed first. We let $j_0(\overrightarrow{e})$ be the jth object of ${}_0(\overrightarrow{e})$. We start counting from 0, so that $0_0(\overrightarrow{e})$ is the head endpoint of e. We choose our ordering to be Λ-invariant: If $\gamma(\overrightarrow{e}) = \overrightarrow{f}$ for some $\gamma \in \Lambda$, then $\gamma(j_0(\overrightarrow{e})) = j_0(\overrightarrow{f})$ for all j. This is possible to arrange because Λ acts freely and transitively on the elements of $\overrightarrow{E}$.

We order the geometric objects of ${}_1(\overrightarrow{e})$ by the rule

$$j_1(\overrightarrow{e}) := j_0(\overrightarrow{e_1}) \qquad \forall j. \tag{7}$$

Let $(\overrightarrow{e})$ be the set of all the geoemtric objects associated to $\overrightarrow{e}$. We order this set by the rule $j(\overrightarrow{e}) = j_0(\overrightarrow{e})$ if $j \leq s$ and $j(\overrightarrow{e}) = (j-s)_1(\overrightarrow{e})$ if $j \geq s+1$. In other words, we list out all the elements of ${}_0(\overrightarrow{e})$ in order and then list out all the elements of ${}_1(\overrightarrow{e})$ in order.

We define V_s as in Step 2 of §1. Let $\overleftarrow{e}$ be the same geodesic as $\overrightarrow{e}$, but with the opposite orientation. The same geometric objects are associated to $\overleftarrow{e}$ but they are listed in a different order. Thus, there is a map $I : V_s \to V_s$ which has the following property:

$$m_j(\overleftarrow{e}) = I(m_j)(\overrightarrow{e}). \tag{8}$$

Note that $I(0_0) = 0_1$. By symmetry our definition is independent of the choice of $\overrightarrow{e}$. Since $\widetilde{\Gamma}$ is a modular pattern there is an involution which preserves $\widetilde{\Gamma}$ and swaps $\overrightarrow{e}$ and $\overleftarrow{e}$. It follows from this structure that I is an involution.

4.3 The Embedding

We make the exact construction we made in Steps 3 and 4 of §1. The only difference is that we use the involution I above, in place of the one defined in the introduction. The rest of the section is devoted to proving Theorem 1.2, with reference to the involution I defined above. We call this the *abstract version* of Theorem 1.2.

Let $\overrightarrow{e_1}$, $\overrightarrow{e_2}$, and $\overrightarrow{e_3}$ by the three geodesic sides of F, the initial ideal triangle in the modular tiling. We orient these geodesics so that they travel counterclockwise in a circuit around F. There is a natural correspondence between the vertices of the simplex Δ_k and the objects $(\overrightarrow{e_k})$. The vertex labelled m_j corresponds to $m_j(\overrightarrow{e_k})$. Each vertex of T belongs to two simplices, but the two labels it gets correspond to the same geometric object. This follows from Equation 7. Thus, the vertices of T naturally correspond to the geometric objects associated to the three edges of F.

The same observations hold for the vertices of T'. In fact, for each counterclockwise oriented triangle of the modular tiling, other than F, we associate a copy of T'. The vertices of T' naturally correspond to the geometric objects associated to the edges of this triangle. When we build the sequence $\{T_m\}$ we are really just gluing together all these copies of T' (along with the single copy of T.)

Here is the key point. If an edge is contained in two adjacent counterclockwise oriented triangles, then it gets opposite orientations from its inclusion into each of the two triangles. For this reason, we need to glue the different copies of T' (and T) together using the map I, which exactly adjusts for the orientation reversal. Thus, our construction of $\{T_m\}$ sets up a correspondence which satisfies the first two conditions of the Correspondence Lemma.

Let $\Lambda' \subset \Lambda$ denote the index 2 subgroup of Λ which does not contain the order 2 elements. If we consider Λ' rather than Λ then we can consistently orient the edges of E so that no element of Λ' reverses the orientations. Namely, we orient the edges of F_0 to make a counterclockwise circuit and then we extend by the action of Λ'. Once we do this, our construction satisfies Condition 3 of the Correspondence Lemma for all $\gamma \in \Lambda'$.

The argument given in §3.4 works here word for word, so that $X(I) = \Delta_\infty$, the limit object associated to a correspondence which satisfies the Correspondence Lemma. This proves that the homeomorphism $Q \to X$ intertwines Λ' with a group of PL maps of X. Finally, we note that there is also a PL map which takes the two simplices of T_1, which are contained in some simplex of T_0, to the two remaining simplices of T_0. When we adjoin this additional PL map we see that our homeomorphism intertwines the full modular symmetry of Q with the PL symmetry of X. This proves X is a PL embedding of Q.

Let X_r and X_i respectively denote the sets of rational and irrational vertices of X, as defined in Step 4 of §1. Since $Q(\Gamma)$ is a topological quotient of the circle, the homeomorphism from $Q(\widetilde{\Gamma})$ to $X(I)$ is induced by a continuous surjection $h : S^1 \to X$. We now explain how h interacts with X_r and X_i.

Lemma 4.1 *X_r and X_i are each dense subsets of X.*

Proof: X consists of those points which are contained in T_m for all m. Any vertex of a simplex in T_m is also a vertex of a simplex in T_{m+1}. Hence every vertex of T_m is contained in X. Hence $X_r, X_i \subset X$. Let $\epsilon > 0$ be given. For sufficiently large m every simplex in T_m has diameter at most ϵ. Since every such simplex has a rational vertex and an irrational vertex, we see that both X_r and X_i are ϵ dense in X. But ϵ is arbitrary. ♠

In §4.1 we constructed an exhaustion $\{F_m\}$ of the ideal triangles of the modular tiling. First of all h maps the three vertices of the triangle F_0 to the three rational vertices of the simplices in T_0. In general, suppose that h maps the vertices of the triangles in F_m to the rational vertices of simplices in T_m. Let e be one of the free edges of F_m. Let Δ_1' and Δ_2' be the two simplices of T_{m+1} contained in $\Delta(e)$. Let v be the vertex of the modular tiling so that the two endpoints of e, together with v, comprise an ideal triangle of F_{m+1}. It follows from the Correspondence Lemma that h maps v to the rational vertex common to Δ_1' and Δ_2'. All in all, h maps the modular group cusps bijectively onto X_r.

Going back to the labelling scheme of §2, each vertex $v \in X_i$ is labelled by a geodesic γ_v. As one travels along γ_v towards an endpoints, one crosses an infinite sequence $\{e_j\}$ of edges of E. The endpoints of e_j are mapped to the rational vertices of a simplex Δ_j such that $v \in \Delta_j$. The nested intersection $\bigcap \Delta_j$ is exactly $\{v\}$. Thus, h maps both endpoints of γ_v to v. If γ is geodesic of $\widetilde{\Gamma}$ then there is some vertex $v \in X_i$ such that $\gamma = \gamma_v$. So, the argument we have just given accounts for all the points in X_i and all the geodesics in $\widetilde{\Gamma}$.

This completes the proof of the abstract version of Theorem 1.2.

5 Combinatorics of the Involution

5.1 Normal Forms

We continue the notation from §4. The goal of this section is to explain how the involution I, constructed in §4.2, depends on the chosen ordering of the elements of ${}_0(\overrightarrow{e})$, and to produce an ordering-independent normal form for I. Lemma 5.1 below connects this normal form to continued fractions.

Let ρ be a permutation of $\{0,...,s\}$, normalized so that $\rho(0)=0$. We define $2\rho : V_s \to V_s$ by the rule $2\rho(m_j) = (\rho(m))_j$. We say that two involutions of V_s are *equivalent* if they are conjugate to each other by a permutation of the form 2ρ. If we choose a different order for ${}_0(\overrightarrow{e})$ there is some permutation ρ of $\{0,...,s\}$ which converts between the two orderings. By Equation 7 and symmetry ρ also converts between the two orderings of ${}_1(\overrightarrow{e})$. From this it follows that the involution I' associated to the new ordering is conjugate to I via 2ρ. In short I and I' are equivalent.

To get a normal form for I we create a decorated graph G_I. There are s vertices of G_I, and these are labelled $1,...,s$. If $i \neq j$ we connect i to j by an edge if $I(i_a) = j_b$ for some choices of $a,b \in \{0,1\}$. We label this edge with 2 arrows, to indicate which of the four cases arises. The general rule is that the arrow near to i points away from i if $i = i_0$ and towards it if $i = i_1$. Likewise, the arrow near to j points away from j if $j = j_0$ and towards it if $j = j_1$.

It cannot happen that $I(i_0) = i_1$. (See the remark after the proof of Lemma 5.1.) If $I(i_0) = i_0$ we adjoin a half-edge to vertex i, which points away from i. If $I(i_1) = i_1$ then we adjoin a half-edge to vertex i, which points to i. Each vertex of G_I has one outgoing edge and one ingoing edge. From this structure we see that each connected component of G_I is a either a decorated circuit or a decorated path. Figure 5.1 below shows an example of each kind.

We call an involution of V_s *normalized* if it pairs 0_0 to 0_1. Two normalized involutions are equivalent iff their graphs are isomorphic. The *normal form* for a normalized involution I is the unordered list of connected components of I, with the labels on the vertices removed.

To each vertex in the normal form we can associate an infinite binary string. If v_0 is a vertex of a circuit, we travel around the circuit endlessly, in the direction indicated by the arrows at v_0. We encounter a seqence of edges $e_1, e_2, e_3, ...$ Let $b_0 = 0$. In general, if $b_k \in \{0,1\}$ we let $b_{k+1} = b_k$ if the two arrows on e_k point in the same direction and $b_{k+1} = 1 - b_k$ otherwise. Our string is then $b_1 b_2 b_3 ...$ For instance, if we start at the top vertex of the circuit on the left hand side of Figure 5.1, we produce the string 10111010... Figure 5.1 shows the scratch-work that produces the initial part of the sequence.

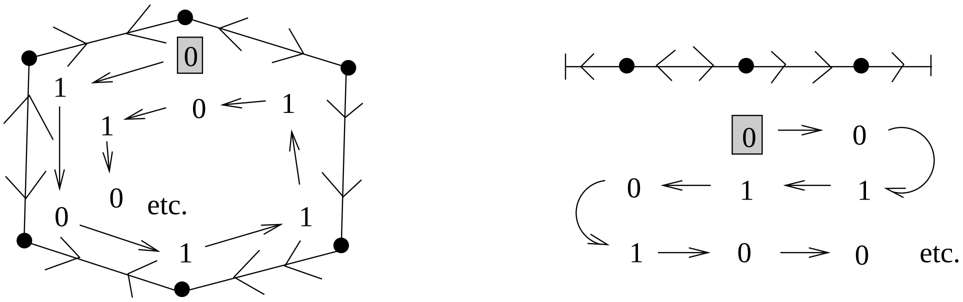

Figure 5.1

If v_0 is a vertex of a path, we start moving in the direction indicated by the arrows at v_0. When we reach the end of a half edge we think of it as a reflector, and bounce back and traverse the half-edge backwards. (This makes it a full edge, on which the two arrows necessarily point in opposite directions.) Thus we travel endlessly back and forth over the path, encountering edges $e_1, e_2, e_3 ...$ We use the same rule as above for generating the binary string. For instance, if v_0 is the middle vertex of the path shown in Figure 5.1 we produce the string 0110100...

Thus, each vertex in our normal form gives rise to an infinite, periodic, binary string. There are exactly s such strings. We call this the *normal form list*.

5.2 Continued Fractions and Cutting Sequences

If $\alpha \in (0,1)$ is a real number, there is a sequence $[a_1 : a_2 ...]$ such that

$$\alpha = \frac{1}{a_1 + \frac{1}{a_2 + ...}} \tag{9}$$

See for instance [**HW**, §10] for details. This sequence is called the *continued fraction expansion* of α. We can convert the sequence $[a_1 : a_2 ...]$ into a binary sequence by putting down a_1 0's, then a_2 1's, then a_3 0's, etc. For instance $[1 : 2 : 3 : 4 ...]$ converts to 0110001111... Following [**S**] we call this the *cutting sequence* of α.

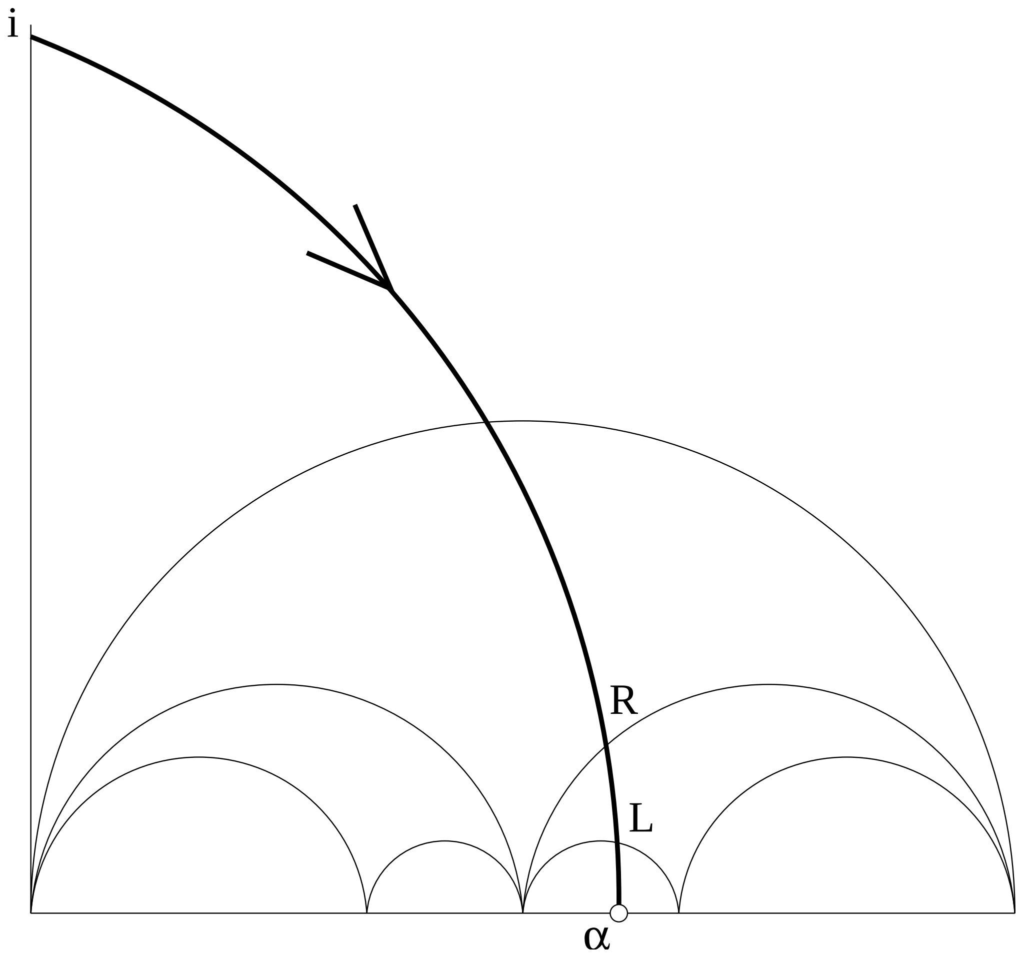

Figure 5.2

We now discuss the content of [**S**, Th. 5.18], which explains the connection between continued fractions and cutting sequences. Thinking of $\boldsymbol{H}^2$ as the upper half plane, we consider the directed geodesic γ which starts at the point i the goes to the point α, as shown in Figure 5.2. As we travel along γ we enter a series of ideal triangles. In passing from one triangle to the next the geodesic either takes a right or a left. Thus, to α we associate a series such as RL... We convert it to binary using the rule $R \to 0$ and $L \to 1$. This procedure recovers the binary continued fraction expansion of α. If α is the endpoint of a geodesic which is the axis of some element of the modular group, the continued fraction expansion of α is periodic, and α is a quadratic irrational.

5.3 Normal Forms and Continued Fractions

We work in the upper half-plane model of $\boldsymbol{H}^2$. Let $\overrightarrow{e}$ be the geodesic connecting ∞ to 0, and pointing towards 0. Let $\overrightarrow{e_0}$ and $\overrightarrow{e_1}$ be the two geodesics defined in §4.2. In our set-up, $\overrightarrow{e_0}$ is the geodesic connecting 0 to 1, and $\overrightarrow{e_1}$ is the geodesic connecting 1 to ∞.

The set ${}_0(\overrightarrow{e_0})$ consists of the geodesics of $\widetilde{\Gamma}$ which cross $\overrightarrow{e}$ and whose head endpoints lie in $(0,1)$. Thus, to $\widetilde{\Gamma}$ we can associate the list of s binary continued fraction expansions of these head endpoints. We call this the *cutting sequence list*. We have already associated another list of s infinite binary strings to $\widetilde{\Gamma}$, namely the *normal form list*. Here is the punchline:

Lemma 5.1 *The cutting sequence list coincides with the normal form list.*

Proof: We begin by giving a slightly different presentation of the geometric construction given in the last section. For $j=0,1$ there is a unique element $\lambda_j \in \Lambda$ such that $\lambda_j(\overrightarrow{e_j}) = \overrightarrow{e}$. Let γ_0 be a geodesic in ${}_0(\overrightarrow{e_0})$. Assume that γ_k has been defined and is a geodesic which crosses $\overrightarrow{e}$. If γ_k crosses $\overrightarrow{e_j}$ we let $\gamma_{k+1} = \lambda_j(\gamma_k)$. The sequence $\{\gamma_k\}$ defines an infinite binary string in the obvious way. The kth digit of this string is the index j of the geodesic $\overrightarrow{e_j}$ crossed by γ_k. It is not hard to see that the string we produce is again the binary continued fraction expansion of the head endpoint.

Let ρ be the order 3 element of Γ which has the orbit $\overrightarrow{e} \to \overrightarrow{e_0} \to \overrightarrow{e_1}$. Let ι be the involution of $\overrightarrow{e}$. We have $\lambda_0 = \iota \circ \rho^2$ and $\lambda_1 = \iota \circ \rho$.

Fix k for the moment and set $\gamma = \gamma_k$. Suppose first that γ crosses $\overrightarrow{e_0}$. There is some m such that $\gamma = m_0(\overrightarrow{e})$. It follows from Equation 7 that

$$\rho^2(\gamma) = m_1(\overrightarrow{e}). \tag{10}$$

By definition,

$$I(m_1)(\overleftarrow{e}) = m_1(\overrightarrow{e}).$$

By symmetry,

$$\iota(I(m_1))(\overleftarrow{e})) = I(m_1)(\overrightarrow{e}).$$

Thus, $I(m_1)(\overrightarrow{e}) = \iota \circ \rho^2(\gamma) = \lambda_0(\gamma) = \gamma_{k+1}$. In short,

$$\gamma_k = m_0(\overrightarrow{e}) \qquad \Rightarrow \qquad \gamma_{k+1} = I(m_1)(\overrightarrow{e}). \tag{11}$$

If γ crosses $\overrightarrow{e}_1$ then the same argument, with ρ in place of ρ^2, shows

$$\gamma_k = m_1(\overrightarrow{e}) \qquad \Rightarrow \qquad \gamma_{k+1} = I(m_0)(\overrightarrow{e}). \tag{12}$$

Equations 11 and 12 say that the sequence $\{\gamma_k\}$ generates elements of V_s by the following two step procedure: First, change the index of the label, so that $m_0 \to m_1$ and $m_1 \to m_0$. Second, apply the involution I.

Recalling the construction of the graph G_I, we see that the above two step procedure generates the succession of vertices we would get by endlessly tracing out the connected component of G_I which contains the number labelling γ_0. We start the tracing with the arrows, because γ_1 belongs to ${}_0(\overrightarrow{e})$. In fact, it is not hard to see that our rule for converting this path into binary tells us exactly whether γ_k lies in ${}_0(\overrightarrow{e}_0)$ or in ${}_1(\overrightarrow{e}_0)$. In other words, the associated normal form string equals the continued fraction string.

Each continued fraction string is a normal form string. All the continued fraction strings are distinct, since two geodesics in $\widetilde{\Gamma}$ cannot have the same endpoint. The two sets of strings coincide, by the pidgeonhole principle. ♠

Remark: If $I(i_0) = i_1$ for some i then, according to Equation 10, the geodesic γ representing i_0 would be stabilized by the parabolic element $\iota \circ \rho^2$. This is impossible.

5.4 Proof of Theorem 1.2

The only missing step in our proof of Theorem 1.2 is the identification of I with the involution defined in Step 2 of §1.

Before we begin our proof we mention some properties of the cutting sequence of a geodesic Γ in a modular pattern. We say that the *minimal sequence* of γ is the minimal finite string which generates its cutting sequence. If γ projects to a loop on the modular orbifold then its minimal sequence has even length. Compare [**S**, Th. 5.16].

We call a finite binary sequence *anti-reflexive* if it has the form $b_1, ..., b_k, a_k, ..., a_1$ where $b_j = 1 - a_j$ for all j. For instance 011001 is reflexive. Now, γ contains an order two fixed point of Λ iff its miminal sequence is anti-reflexive. So, if γ is the lift of a *smooth* loop on the modular orbifold then its minimal sequence is not anti-reflexive.

If $\widetilde{\Gamma}$ is a minimal modular pattern then the cutting sequence list consists of a single string, and it is the string $\{b_k\}$ constructed in Step 1 of §1. Hence the normal form list is a single binary string, and G_I has only one connected component. If this connected component was a segment, then the minimal sequence for each γ in $\widetilde{\Gamma}$ would be anti-reflexive. This is not the case. Hence the one component of G_I is a cycle.

The rules we give in Step 2 of §1 are nothing more than the method which converts a cycle of G_I, together with a chosen vertex, into a binary string on the normal form list. Thus, in the case of minimal modular patterns, the involution I defined in §4.2 coincides with the one defined in Step 2 of §1. This observation, together with the work in §4.3, completes the proof of Theorem 1.2.

Remark: As we promised in the introduction we explain the modification we need to make when $\widetilde{\Gamma}$ is the lift of a closed loop on the modular orbifold which does contain the order 2 point. In this case, the minimal sequence for any geodesic γ in the pattern is anti-reflexive, and has the form $b_1, ..., b_k, a_k, ..., a_1$, as above. In this case, the sequence we use in Step 1 is $b_1, ..., b_k$. The involution we create in Step 2 is defined by the rule:

$$I(k_{b_k}) = (k+1)_{1-b_{k+1}} \qquad s = 1, ..., k-1;$$
$$I(1_{1-b_1}) = 1_{1-b_1};$$
$$I(s_{b_s}) = s_{b_s}. \tag{13}$$

For instance, let $\widetilde{\Gamma}$ be the modular pattern generated by a geodesic γ which joins the midpoints of two sides of an ideal triangle in the modular tiling. (The midpoints are the fixed points of the relevant order 2 elements of Λ.) The endpoints of γ lie in $\boldsymbol{Q}[\sqrt{5}]$. The cutting sequence for γ has the form ...010101... The minimal sequence, 01, is anti-reflexive. The sequence used in Step 1 is 0, and $s = 1$. The involution I in Step 2 has the action

$$0_0 \leftrightarrow 0_1; \quad 1_1 \to 1_1 \quad 1_0 \to 1_0.$$

This involution produces the three dimensional embedding of [**S2**], which is topologically the same as the embedding of the *prototypical example* in [**S1**].

6 References

[**HW**] G.H. Hardy and E.M. Wright, *An Introduction to the Theory of Numbers*, 4th edition, Oxford University Press, 1959.

[**S**] C. Series, *Geometric Methods of Symbolic Coding*, in Ergodic Theory, Symbolic Dynamics, and Hyperbolic Spaces (ed. Bedford, Keane and Series.) Oxford Science Publications (1991)

[**S1**] R.E. Schwartz, *Circle Quotients and String Art*, Topology (2001)

[**S2**] R.E. Schwartz, *Degenerating the Complex Hyperbolic Ideal Triangle Groups*, Acta Mathematica **186** (2001)

[**T**] W. Thurston, *Three Dimensional Manifolds, Kleinian Groups, and Hyperbolic Geometry*, Bulletin of the A.M.S., 1983.