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

Geometry & Topology
Volume 8 (2004) 1–34
Published: 18 January 2004

Modular circle quotients and PL limit sets

Richard Evan Schwartz

Department of Mathematics, University of Maryland
College Park, MD 20742, USA

Email: `res@math.umd.edu`

Abstract

We say that a collection Γ of geodesics in the hyperbolic plane $\boldsymbol{H}^2$ is a *modular pattern* if Γ is invariant under the modular group $PSL_2(\boldsymbol{Z})$, if there are only finitely many $PSL_2(\boldsymbol{Z})$–equivalence classes of geodesics in Γ, and if each geodesic in Γ is stabilized by an infinite order subgroup of $PSL_2(\boldsymbol{Z})$. For instance, any finite union of closed geodesics on the modular orbifold $\boldsymbol{H}^2/PSL_2(\boldsymbol{Z})$ lifts to a modular pattern. Let S^1 be the ideal boundary of $\boldsymbol{H}^2$. Given two points $p,q \in S^1$ we write $p \sim q$ if p and q are the endpoints of a geodesic in Γ. (In particular $p \sim p$.) We will see in section 3.2 that $\sim$ is an equivalence relation. We let $Q_\Gamma = S^1/\sim$ be the quotient space. We call Q_Γ a *modular circle quotient*. In this paper we will give a sense of what modular circle quotients "look like" by realizing them as limit sets of piecewise-linear group actions

AMS Classification numbers Primary: 57S30

Secondary: 54E99, 51M15

Keywords: Modular group, geodesic patterns, limit sets, representations

Proposed: David Gabai
Seconded: Martin Bridson, Walter Neumann

Received: 4 February 2003
Accepted: 13 January 2004

1 Introduction

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? Rather than consider the expensive disasters produced by an actual experiment we will consider related theoretical objects called *modular circle quotients*.

We say that a collection Γ of geodesics in the hyperbolic plane $\boldsymbol{H}^2$ is a *modular pattern* if Γ is invariant under the modular group $PSL_2(\boldsymbol{Z})$, if there are only finitely many $PSL_2(\boldsymbol{Z})$–equivalence classes of geodesics in Γ, and if each geodesic in Γ is stabilized by an infinite order subgroup of $PSL_2(\boldsymbol{Z})$. For instance, any finite union of closed geodesics on the modular orbifold $\boldsymbol{H}^2/PSL_2(\boldsymbol{Z})$ lifts to a modular pattern. Let S^1 be the ideal boundary of $\boldsymbol{H}^2$. Given two points $p, q \in S^1$ we write $p \sim q$ if p and q are the endpoints of a geodesic in Γ. (In particular $p \sim p$.) We will see in section 3.2 that $\sim$ is an equivalence relation. We let $Q_\Gamma = S^1/\sim$ be the quotient space. We call Q_Γ a *modular circle quotient*.

In [7] we encountered a certain modular circle quotient as the limit set of a special representation of $PSL_2(\boldsymbol{Z})$ into $PU(2,1)$, the group of complex projective automorphisms of the 3–sphere S^3. In [8] we embedded some related circle quotients into S^3. In this paper we will treat all the modular circle quotients, motivated by the constructions in [8] but starting from scratch. Our aim is to give a sense of what they look like, by realizing them as limit sets of piecewise linear group actions.

1.1 Statement of results

Let Γ be a modular pattern of geodesics. As we explain in section 3.1, there is a well-known tiling of $\boldsymbol{H}^2$ by ideal triangles which is invariant under the action of $PSL_2(\boldsymbol{Z})$. We call this tiling the *modular tiling*. We define $|\Gamma|$ to be one more than the number of geodesics in Γ which intersect a given edge of the modular tiling. We will see in section 3.2 that this number is finite. $|\Gamma|$ is independent of the choice of edge, by symmetry.

Let S^n be the n–sphere. Our model for S^n is the double of an n–simplex: $S^n = \Delta_+ \cup \Delta_-$, where Δ_+ and Δ_- are two copies of an n–simplex, glued along their boundaries. A *simplex* of S^n is a sub-simplex of either Δ_+ or Δ_-. Say that a *punctured simplex* of S^n is a simplex with its vertices deleted.

A homeomorphism h of S^n is *piecewise linear* (or PL) if there is some triangulation of S^n into finitely many simplices such that h is affine when restricted to each simplex in the triangulation. The set of PL homeomorphisms of S^n forms a topological group $\mathrm{PL}(S^n)$ equipped with the compact–open topology. Let $H \subset \mathrm{PL}(S^n)$ be a subgroup. A compact subset $\Lambda \subset S^n$ is the *limit set* of H if H acts *properly discontinuously* on $S^n - \Lambda$ and *minimally* on Λ. Thus $H(x)$ is dense in Λ for every $x \in \Lambda$ and for any compact $K \subset S^n - \Lambda$, the set $\{h \in H \mid h(K) \cap K \neq \emptyset\}$ is finite.

Theorem 1.1 *Let $n = |\Gamma|$. There is an embedding $i\colon Q_\Gamma \to S^n$ and a monomorphism $\rho\colon PSL_2(\boldsymbol{Z}) \to \mathrm{PL}(S^n)$ such that $i(Q_\Gamma)$ is the limit set of $\rho(PSL_2(\boldsymbol{Z}))$. There is a $\rho(PSL_2(\boldsymbol{Z}))$–invariant partition of $S^n - i(Q_\Gamma)$ into punctured simplices, the vertices of which are densely contained in $i(Q_\Gamma)$.*

One generalization of a modular pattern is a $PSL_2(\boldsymbol{Z})$–invariant map $f\colon \Gamma \to (0,1]$, where Γ is a modular pattern. Let $\Box(\Gamma)$ be the space of these maps. Let $CS(S^n)$ be the space of closed subsets of S^n, given the Hausdorff topology. (Two subsets are close if each is contained in a small tubular neighborhood of the other.) Let $\mathrm{Mon}(PSL_2(\boldsymbol{Z}), \mathrm{PL}(S^n))$ denote the space of monomorphisms from $PSL_2(\boldsymbol{Z})$ into $\mathrm{PL}(S^n)$ given the algebraic topology. (Two monomorphisms are close if they map the generators to nearby elements of $\mathrm{PL}(S^n)$.) The following result organizes all the modular circle quotients based on subpatterns Γ' of Γ.

Theorem 1.2 *Let $n = |\Gamma|$. There are continuous maps $\Lambda\colon \Box(\Gamma) \to CS(S^n)$ and $\rho\colon \Box(\Gamma) \to \mathrm{Mon}(PSL_2(\boldsymbol{Z}), \mathrm{PL}(S^n))$ such that the following is true for all $f \in \Box(\Gamma)$. The set Λ_f is the limit set of $\rho_f(PSL_2(\boldsymbol{Z}))$ and Λ_f is homeomorphic to $Q_{\Gamma'}$, where $\Gamma' = f^{-1}(1)$.*

The method we use to prove Theorem 1.2 is flexible and allows us to make a statement about more general kinds of circle quotients:

Theorem 1.3 *Let Γ' be the lift to H^2 of an arbitrary finite union of closed geodesics on a cusped hyperbolic surface Σ. Let $Q_{\Gamma'}$ be the circle quotient based on Γ'. For some n there is an embedding $i\colon Q_{\Gamma'} \to S^n$ and a monomorphism $\rho\colon \pi_1(\Sigma) \to \mathrm{PL}(S^n)$ such that $i(Q_{\Gamma'})$ is the limit set of $\rho(\pi_1(\Sigma))$.*

If Γ_1 and Γ_2 are both modular patterns and $\Gamma_1 \subset \Gamma_2$ then we have an inclusion $\Box(\Gamma_1) \hookrightarrow \partial\Box(\Gamma_2)$. Assuming this inclusion implicitly, we say that a sequence $\{f_m\} \in \Box(\Gamma_2)$ *degenerates* to $f \in \Box(\Gamma_1)$ if $f_m(\gamma) \to 0$ for all $\gamma \in \Gamma_2 - \Gamma_1$ and $f_m(\gamma) \to f(\gamma)$ if $\gamma \in \Gamma_1$. Let $n_j = |\Gamma_j|$ for $j = 1, 2$. The n_2–simplex has faces which are n_1–simplices. The doubles of these n_1–simplices are copies of S^{n_1} contained in S^{n_2}. We call these copies the *natural embeddings* of S^{n_1} into S^{n_2}.

Theorem 1.4 *There is a natural embedding $i\colon S^{n_1} \hookrightarrow S^{n_2}$ with the following property. Let $\{f_m\} \in \square(\Gamma_2)$ be a sequence which degenerates to $f \in \square(\Gamma_1)$. Then the limit sets Λ_{f_m} converge to $i(\Lambda_f)$. The restriction of ρ_{f_m} to Λ_{f_m} converges to the action of $i \circ \rho_f \circ i^{-1}$ on $i(\Lambda_f)$.*

Theorem 1.4 covers one case not explicitly mentioned. In section 5.5 we define a certain *standard representation* $\rho_0\colon PSL_2(\boldsymbol{Z}) \to PL(S^1)$. If $f_m(\gamma) \to 0$ for all $\gamma \in \Gamma_2$ then Λ_{f_m} converges to a naturally embedded circle $i(S^1)$ and the restriction of ρ_{f_m} to Λ_{f_m} converges to $i \circ \rho_0 \circ i^{-1}$.

Theorem 1.4 lets us organize all the modular circle quotients into a coherent whole. We define $\square = \bigcup_\Gamma \square(\Gamma)$. Let S^∞ be the direct limit of S^n, under the system of natural embeddings. Let $CS(S^\infty)$ denote the set of finite dimensional closed subsets of S^∞ equipped with the Hausdorff topology. Then Theorem 1.4 gives a map $\Lambda\colon \square \to CS(S^\infty)$ such that Λ_f is homeomorphic to $Q_{\Gamma'}$ and contained in a naturally embedded $|\Gamma|$–dimensional sphere. Here $\Gamma = f^{-1}((0,1])$ and $\Gamma' = f^{-1}(1)$. The map Λ is continuous when restricted to each finite dimensional subspace of $\square$.

The following construction illustrates the nature of our results. List all the vertices $v_0, v_1, v_2 ...$ of $\square$ with v_0 being the vertex corresponding to the 0–map—i.e. the empty pattern. Let $\{f_t \mid t \in [0,\infty)\} \subset \square$ be a continuous path such that $f|_{[0,n]}$ is contained in the convex hull of the vertices $v_0, ..., v_n$ and $f_n = v_n$. Here $n = 0, 1, 2....$ Then Λ_0 is just the double of a line segment. As t increases Λ_t continuously and endlessly crinkles up, assuming the topology of every modular circle quotient as it goes.

1.2 Comparisons and speculation

Here are some possible connections to our results:

(1) Our constructions here are similar in spirit to our constructions in [9], where we related the modular group to Pappus's theorem and thereby produced discrete representations of the modular group into the group of automorphisms of the real projective plane.

(2) Our Theorem 1.1 seems at least vaguely related to the general results in [2] about embedding the boundaries of hyperbolic groups into S^n.

(3) Some of the combinatorial ideas underlying our constructions are related to the theory [6] of coding geodesics on the modular surface using their cutting sequences. We can work this out explicitly but don't do it in this paper.

(4) $\square(\Gamma)$ (with its associated maps) is like a PL version of Teichmuller space. The groups attached to the set $\{f \mid f^{-1}(1) = \emptyset\}$ are like PL quasi-Fuchsian groups [1, 4] in that their limit sets are topological circles. The other groups are like cusp groups on the boundary of quasi-Fuchsian space.

We elaborate on the fourth item. $\square(\Gamma)$ is both richer and poorer than Teichmuller space. It is richer because it allows for deformations which cannot exist in hyperbolic geometry. There are no nontrivial deformations of the modular group into $\mathrm{Isom}(\boldsymbol{H}^3)$ whereas $\dim \square(\Gamma)$ grows unboundedly with the complexity of Γ. Indeed, one possible use of our results is that they provide a topological model for degenerating families of representations of punctured surface groups—i.e. finite index subgroups of the modular group—into a Lie group. Such families generally are extremely difficult to construct, let alone study geometrically. Our results give a glimpse of how punctured surface groups might degenerate when *non-simple* closed geodesics on the surface are pinched.

$\square(\Gamma)$ is poorer than Teichmuller space because it only allows for degenerations which occur by pinching closed geodesics. We don't get things like geometrically infinite limits. It almost goes without saying that $\square(\Gamma)$ is geometrically much poorer than Teichmuller space. It does not enjoy any of the beautiful rigid structure [3] of Teichmuller space.

We wonder how our results transfer to the more rigid setting of a Lie group G acting on a homogeneous space X. We think that it ought to be possible sometimes to geometrize our constructions and produce representations of $PSL_2(\boldsymbol{Z})$ into G which "realize" our PL representations. The result in [7] is an example of this. On the other hand, we think that there should be strong restrictions on the types of circle quotients for each pair (G, X). A general *restriction result* would provide a new tool in the study of representations of surface groups into Lie groups, because it would help control the possible degenerations.

As far as we know, all the modular circle quotients are non-planar. At any rate, many of them are non-planar and hence cannot be embedded into S^2. Probably all of the modular circle quotients can be embedded into S^3. However, such embeddings would probably be very "distorted" in general. We would like to quantify this distortion, and relate it to the complexity of the modular pattern.

We also wonder about how our results work out for circle quotients based on uniform lattices, but don't have any idea how to proceed.

1.3 Some ideas in the proof

Our main idea is to construct an object we call a *modular block* (or *block* for short.) A block is a certain subset $\Omega \subset S^n$ equipped with an order 3 PL automorphism σ. A block is based on a neat partition of the n–simplex into $3^k - 1$ smaller n–simplices. Here $k = (n+1)/2$, with n always being odd. The partition is combinatorially isomorphic to the k–fold join of a triangle (which is an n–sphere) minus one n–simplex. Ω is obtained by deleting 2 simplices from the partition, so that $\partial\Omega$ consists of 3 non-disjoint n–simplex boundaries, called *terminals*. The remaining $3^k - 3$ simplices partition Ω and are permuted by σ.

We will construct an infinite network of blocks glued together along terminals. The network is essentially tree-like but its fine structure is related to the symbolic coding of geodesics in Γ. It turns out that Q_Γ is homeomorphic to the closure of the block vertices. $PSL_2(\boldsymbol{Z})$ is represented as a subgroup of the automorphism group of the network. Underlying our block network is a kind of correspondence between some hyperbolic geometry objects related to the modular tiling and some simplicial objects. We call this a *simplicial correspondence*. The following table summarizes the correspondence.

hyperbolic object	**simplicial object**
the modular tiling T	modular block network
ideal triangle of T	modular block
geodesic edge of T	terminal
ideal vertex of T; geodesic of Γ	vertex of a block.
circle quotient	closure of the block vertices

Here is a more global point of view. We can define an abstract simplicial complex $C(\Gamma)$ whose vertices are elements of $\Gamma \cup VT$. Here VT is the set of ideal vertices of the modular tiling. We say that a subset $S \subset \Gamma \cup VT$ is an *abstract simplex* if it satisfies the following properties:

(1) There is some ideal triangle τ of T (not necessarily unique) such that every $s \in S$ is either an ideal vertex of τ or a geodesic of Γ which intersects τ. We say that τ and S are *associated*.

(2) If τ is associated to S and $H_\tau \subset PSL_2(\boldsymbol{Z})$ is the order–3 stabilizer subgroup of τ then S does not contain an orbit of H_τ. Moreover, S is not stabilized by an order 2 element of $PSL_2(\boldsymbol{Z})$.

Evidently $PSL_2(\boldsymbol{Z})$ acts on $C(\Gamma)$. It turns out that the maximal abstract simplices of $C(\Gamma)$ are n–dimensional and that $C(\Gamma)$ minus the vertices is a combinatorial n–manifold. There are 3^k-3 maximal abstract simplices of $C(\Gamma)$ associated to each τ. Our construction gives an embedding of $C(\Gamma)$ into S^n in such a way that these 3^k-3 abstract simplices map to the simplices partitioning the block corresponding to τ. The embedding conjugates the natural action of $PSL_2(\boldsymbol{Z})$ on $C(\Gamma)$ to a subgroup of the automorphism group of the block network. The embedding maps the vertex set of $C(G)$ to a dense subset of the limit set.

So far we have sketched the proof of Theorem 1.1. For the remaining results, our idea is to modify the block network by a certain 2–step process. First, we push the blocks apart from each other by attaching collar-like sets, which we call *separators*, onto the block terminals. Compare Figure 5.1. This process allows the topology of the limit set to vary with the stratum of $\square(\Gamma)$, as in Theorem 1.2. (Theorem 1.3 comes as another application.) Second, we *warp* the shapes of the individual blocks, to allow the representations associated to $\square(\Gamma_2)$ to degenerate to the representations associated to $\square(\Gamma_1)$, as in Theorem 1.4. The element of $\square(\Gamma)$ determines both the shapes of the warped blocks and the shapes of the separators.

1.4 Overview of the paper

We have tried to make this paper completely self-contained. It only relies on a few basic ideas from linear algebra, hyperbolic geometry, and real analysis. We remark to the interested reader that section 2 and 3 makes for a complete, shorter paper in itself, which proves Theorem 1.1. Here is a plan of the rest of paper:

Section 2: Modular blocks, containing: 2.1: The Block Lemma; 2.2: The details; 2.3: 3–dimensional example.

Section 3: Theorem 1.1, containing: 3.1: The modular tiling; 3.2: Modular pattern basics; 3.3: Simplicial correspondences; 3.4: Embedding the quotient; 3.5: Block networks; 3.6: Putting it together.

Section 4: Modified blocks, containing: 4.1: Partial prisms; 4.2: Separators; 4.3: Warped blocks; 4.4: Main construction; 4.5: Degeneration.

Section 5: The rest of the results, containing: 5.1: Modified correspondences; 5.2: Modified block networks; 5.3: Proof of Theorem 1.2; 5.4: Proof of Theorem 1.3; 5.5: Proof of Theorem 1.4.

References

Acknowledgements I would like to thank the IHES for their hospitality during the writing of an early version of this paper.

The author is supported by NSF Research Grant DMS-0072607.

2 Modular blocks

2.1 The Block Lemma

Let $k \geq 2$ be an integer and let $n = 2k - 1$. Let Δ_0 be an n–simplex. We say that a *modular block* is a set

$$\Omega = \text{closure}(\Delta_0 - \Delta_1 - \Delta_2) \tag{1}$$

Where $\Delta_1, \Delta_2 \subset \Delta_0$ are n–simplices with disjoint interiors and

(1) For any indices $i \neq j$ there are k vertices common to Δ_i and Δ_j, and $\partial\Delta_i \cap \partial\Delta_j$ is the convex hull of these common vertices.

(2) There is an order 3 PL automorphism $\sigma\colon \Omega \to \Omega$ such that σ is affine on $\partial\Delta_j$, with orbit $\partial\Delta_0 \to \partial\Delta_1 \to \partial\Delta_2 \to \partial\Delta_0$.

We call $\partial\Delta_j$ a *terminal* of Ω for $j = 0, 1, 2$. We call $\partial\Delta_0$ the *outer terminal* and $\partial\Delta_1$ and $\partial\Delta_2$ the *inner terminals*.

Recall from section 1.1 that $S^n = \Delta_+ \cup \Delta_-$, where Δ_+ and Δ_- are two copies of a standard n–simplex. Our model for $\Delta_\pm$ is the convex hull of the standard basis vectors in $\boldsymbol{R}^{n+1} = \boldsymbol{R}^{2k}$. The goal of this chapter is to prove

Lemma 2.1 (Block Lemma) *There exists a modular block whose outer terminal is $\partial\Delta_+$.*

Proof – modulo some details Let $e_1, ..., e_k$ be the standard basis vectors in $\boldsymbol{R}^k$. For any $r \in \boldsymbol{R}$, let $r_{(k)} = (r, ..., r) \in \boldsymbol{R}^k$. For $j = 1, ..., k$ we define the following points of $\Delta_0 = \Delta_+$:

$$A_j = (e_j, 0_{(k)}); \qquad B_j = \frac{1}{2n}(2_{(k)} - e_j, 2_{(k)} - e_j); \qquad C_j = (0_k, e_j); \tag{2}$$

Let $Y = \{Y_j\}_{j=1}^k$ for each letter $Y \in \{A, B, C\}$. Let $\langle\cdot\rangle$ denote the convex hull operation. Note that $\Delta_0 = \langle A \cup C\rangle$. We define

$$\Delta_1 = \langle A \cup B\rangle; \qquad \Delta_2 = \langle B \cup C\rangle. \tag{3}$$

The sets $A \cup B$ and $B \cup C$ are bases for $\boldsymbol{R}^{2k}$. (See Lemma 2.2.) Hence Δ_1 and Δ_2 are n–simplices. Define $u = (1_{(k)}, -1_{(k)})$. We have $B_j \cdot u = 0$ for all j. Therefore B is contained in the hyperplane $u^\perp$. We also have $A_j \cdot u = 1$ and $C_j \cdot u = -1$ for all j. Therefore $u^\perp$ separates A from C. Hence $\Delta_1 \cap \Delta_2 = \langle B \rangle$. Since $B \in \mathrm{int}(\Delta_0)$ we have $\partial\Delta_0 \cap \partial\Delta_1 = \langle A \rangle$. Likewise $\partial\Delta_0 \cap \partial\Delta_2 = \langle C \rangle$. Thus Δ_0, Δ_1, and Δ_2 satisfy Condition 1.

Let $X = A \cup B \cup C$. We define $\sigma\colon X \to X$ by the action

$$A_j \to B_j \to C_j \to A_j; \qquad j = 1, ..., k. \tag{4}$$

Equation 3 implies that σ extends to a self-homeomorphism of $\partial\Omega$, which is affine on each terminal. Here Ω is as in Equation 1.

Say that a $2k$–element $S \subset X$ is *good* if it does not contain any orbits of σ and does not equal $A \cup C$. There are $3^k - 1$ good subsets, two of which are $A \cup B$ and $B \cup C$. Lemma 2.2 below shows that every good set is a basis. We define a *good simplex* to be a simplex of the form $\langle S \rangle$, where S is a good subset. We can extend the action of σ to any individual good simplex other than Δ_1 and Δ_2 by the rule $\sigma(\langle S \rangle) = \langle \sigma(S) \rangle$.

We will show below that Δ_0 is triangulated by the good simplices. That is, $\Delta_0 = \bigcup_S \langle S \rangle$, and for all good simplices $\langle S_1 \rangle$ and $\langle S_2 \rangle$, we have

$$\langle S_1 \rangle \cap \langle S_2 \rangle = \langle S_1 \cap S_2 \rangle. \tag{5}$$

Ω is triangulated by the good simplices which are not Δ_1 or Δ_2, and these are permuted by σ. Equation 5 implies that all the individual actions of σ on good simplices fit together continuously. Hence $\sigma\colon \Omega \to \Omega$ satisfies Condition 2. $\square$

2.2 The details

For $b \geq 1$ we introduce the $b \times b$ matrix Υ_b whose (ij)th entry is 1 if $i = j$ and otherwise 2. This circulent matrix has the eigenvalue $2b - 1$ with multiplicity 1 and the eigenvalue -1 with multiplicity $b - 1$. Therefore

$$(-1)^{b-1} \det(\Upsilon_b) > 0. \tag{6}$$

Before proving Lemma 2.2 let's consider a representative example which shows how Υ_b arises in our calculations. We take $(k, n) = (3, 5)$ and show that the set $S = \{A_1, B_1, B_2, C_2, A_3, C_3\}$ is a basis for $\boldsymbol{R}^6$. Let M be the matrix whose rows are elements of S. If some row has a single 1 in the jth spot, and 0s in all other spots, we change jth spots of all the other rows to 0. We call this *simple row reduction*. We use a combination of permutations and simple row reductions to show that $\det(M) \neq 0$. Ignoring the factor of $\frac{1}{2n}$ in the second and third rows:

$$\begin{bmatrix} \underline{1} & 0 & 0 & 0 & 0 & 0 \\ 1 & 2 & 2 & 1 & 2 & 2 \\ 2 & 1 & 2 & 2 & 1 & 2 \\ 0 & 0 & 0 & 0 & \underline{1} & 0 \\ 0 & 0 & \underline{1} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \underline{1} \end{bmatrix} \to \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \to \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 2 \\ 0 & 0 & 0 & 0 & 2 & 1 \end{bmatrix}$$

This last matrix obviously has nonzero determinant. Notice also that Υ_2 appears in the bottom right corner, and 2 is the cardinality of $S \cap B$.

Lemma 2.2 *Every good set is a basis of* $\boldsymbol{R}^{2k}$.

Proof Let S be a good set. Let b be the cardinality of $S \cap B$. Using permutations and simple row reduction we see that

$$\det(M) = s \det \begin{bmatrix} I & 0 \\ 0 & \frac{1}{2n}\Upsilon_b \end{bmatrix} \neq 0. \tag{7}$$

Here $s \in \{-1, 1\}$ depends on the number of permutations. □

Lemma 2.3 *Let $\langle S_1 \rangle$ be a good simplex. Each codimension–1 face $\langle S' \rangle$ of $\langle S_1 \rangle$, which is not a face of Δ_0, is a face of one other good simplex $\langle S_2 \rangle$. Equation 5 holds for $\langle S_1 \rangle$ and $\langle S_2 \rangle$.*

Proof We have $S' = S_1 - Y_j$ for some $j \in \{1, ..., k\}$ and $Y_j \in \{A_j, B_j, C_j\}$. Without loss of generality assume $j = 1$. By hypotheses $S' \not\subset A \cup C$. Hence there is exactly one other way to complete S' to a good subset: Namely, $S_2 = S' \cup Z_1$, where $Z_1 = \{A_1, B_1, C_1\} - S_1$. Let M_Y and M_Z denote the matrices whose rows are the elements of S_1 and S_2 respectively. We require that Y_1 and Z_1 appear in the same rows of M_Y and M_Z respectively and that all other rows coincide. To verify Equation 5 for $\langle S_1 \rangle$ and $\langle S_2 \rangle$ it suffices to prove that $\det(M_Y)/\det(M_Z) < 0$. The idea here is that this causes Y_1 and Z_1 to lie on opposite sides of the hyperplane containing $\langle S' \rangle$. By symmetry it suffices to consider the cases $(Y, Z) = (B, C)$ and $(Y, Z) = (A, C)$. We will consider these in turn.

Case 1 Let b be the cardinality of $S \cap B$. Since $S' = S_1 - B_1 \not\subset A \cup C$ we have $b \geq 2$. Using the operations of Lemma 2.2 we get the formula in Equation 7 for $\det(M_B)$. When we perform the same operations on M_C we get the same matrix as in Equation 7, except that all the 2's in one of the rows are changed to 0's. We can then perform one more simple row reduction, using this row, to get

$$\det(M_C) = s \det \begin{bmatrix} I & 0 \\ 0 & \frac{1}{2n} \Upsilon_b^{(a)} \end{bmatrix} \tag{8}$$

for some $a \in \{1, ..., b\}$. Here $\Upsilon_b^{(a)}$ is created from Υ_b by changing the (aj)th and (ja)th entries of Υ_b from 2 to 0, for all $j \neq a$. Independent of a we have

$$\det(\Upsilon_b^{(a)}) = \det(\Upsilon_{b-1}). \tag{9}$$

Equations 6-9 give $\det(M_B)/\det(M_C) < 0$.

Case 2 Suppose that A_1 and B_1 are the first two rows of M_A and that C_1 and B_1 are the first two rows of M_C. Let M be the matrix obtained by replacing the first row of M_A (or M_C) by $A_1 + C_1$. We have $\det(M) = \det(M_A) + \det(M_C)$. The first row of M is $(1, 0, ..., 0, 1, 0, ..., 0)$. Using this row for row-reduction we can make all other rows have zeros in the 1st and $(k+1)$st positions. The last $2k - 2$ rows of M are linearly independent by Lemma 2.2. Therefore we can perform a series of row reductions to change the remaining entries of the second row of M to 0s. Hence $\det(M) = 0$ and $\det(M_A)/\det(M_C) = -1$. □

Remark To be sure we checked all the calculations entailed by the preceding lemma by computer for the cases $n = 3, 5, 7, 9, 11, 13$.

Corollary 2.4 *Δ_0 is the union of the good simplices.*

Proof Let Δ' be the union of the good simplices. Δ' is closed subset of Δ_0. If $\Delta' \neq \Delta_0$ then some codimension one subset of $\partial\Delta'$ separates the nonempty $\mathrm{int}(\Delta_0 - \Delta')$ from the nonempty $\mathrm{int}(\Delta')$. Hence there is a good simplex $\langle S_1 \rangle$, a codimension 1 face $\langle S' \rangle$ of $\langle S_1 \rangle$, and a point $x \in \mathrm{int}(\Delta_0) \cap \mathrm{int}(\langle S' \rangle) \cap \partial\Delta'$. Note that $\langle S' \rangle$ is not a face of Δ_0. By Lemma 2.3 there is a good simplex $\langle S_2 \rangle$ which also has $\langle S' \rangle$ as a face, and $x \in \mathrm{int}(\langle S_1 \rangle \cup \langle S_2 \rangle) \subset \mathrm{int}(\Delta')$. This is a contradiction. □

Corollary 2.5 *Equation 5 is true for all good simplices $\langle S_1 \rangle$ and $\langle S_2 \rangle$ provided that $\langle S_1 \cap S_2 \rangle$ has codimension less than 3.*

Proof Lemma 2.3 takes care of the codimension 1 case. Let $F = \langle S_1 \cap S_2 \rangle$ have codimension 2. We will treat the case when F is not a face of Δ_0, the other case being very similar.

There are either 3 or 4 ways to complete $S_1 \cap S_2$ to a good set. From Lemma 2.3 the corresponding good simplices just wind around F in a cyclic fashion. That is, there is a cyclic ordering to the simplices, such that consecutive simplices are as in Lemma 2.3. The simplices are prevented from winding more than once around F by the fact that the total dihedral angle around F is less than $4 \times \pi = 4\pi$. The topological situation just described implies that $\langle S_1 \rangle \cap \langle S_2 \rangle = F$. □

Corollary 2.6 *Equation 5 holds for every pair of good simplices.*

Proof Let $\Delta = \Delta_0$. Let $\widehat{\Delta}$ be the abstract simplicial complex obtained by gluing together all the good simplices along the convex hulls of their common vertices. We have a tautological map $I \colon \widehat{\Delta} \to \Delta$ which maps each abstract version of a good simplex to its realization as a subset of Δ. It suffices to prove that I is a bijection.

Let $\widehat{\Delta}_k$ denote the interior of the complement of the codimension–k skeleton of $\widehat{\Delta}$. Corollary 2.5 implies that I is a local isometry on $\widehat{\Delta}_3$. The point here is that we just need to look at the links of interior simplices of codimension 1 and 2, and this is what we have done.

I maps the codimension–3 skeleton of $\widehat{\Delta}$ onto the set of codimension–3 faces of the good simplices. Since I is onto (by Corollary 2.4) the set $\Delta_3 = \mathrm{int}(I(\widehat{\Delta}_3))$ is obtained from $\mathrm{int}(\Delta)$ by deleting the codimension–3 faces. Hence Δ_3 is open, simply connected and dense. We can find a local isometry J, defined on an open subset of Δ_3, which is the inverse of I where defined. Since Δ_3 is open and simply connected, J extends by analytic continuation to a local isometry on Δ_3. Since Δ_3 is dense, J extends to all of Δ.

Since $\widehat{\Delta}_3$ and Δ_3 are both simply connected it follows from analytic continuation that the local isometries $I \circ J$ and $J \circ I$ are the identity on Δ_3 and $\widehat{\Delta}_3$ respectively. By continuity, they are the identity on Δ and $\widehat{\Delta}$ respectively. Hence I is a bijection. □

2.3 3–dimensional example

We illustrate our construction by working out the 3 dimensional case more explicitly. To get a 3–dimensional picture we use the projection

$$V \to (V \cdot (1,1,-1,-1), V \cdot (1,-1,1,-1), V \cdot (1,-1,-1,1))$$

Using this projection we have

$$A_1:\begin{bmatrix}1\\1\\1\end{bmatrix}\quad A_2:\begin{bmatrix}1\\-1\\-1\end{bmatrix}\quad B_1:\begin{bmatrix}0\\-\frac{1}{3}\\0\end{bmatrix}\quad B_2:\begin{bmatrix}0\\\frac{1}{3}\\0\end{bmatrix}\quad C_1:\begin{bmatrix}-1\\1\\-1\end{bmatrix}\quad C_2:\begin{bmatrix}-1\\-1\\1\end{bmatrix}$$

Figure 2.1 shows a projection to the xy plane. The two tetrahedra on the right are supposed to fit inside the one on the left, as indicated by the labels.

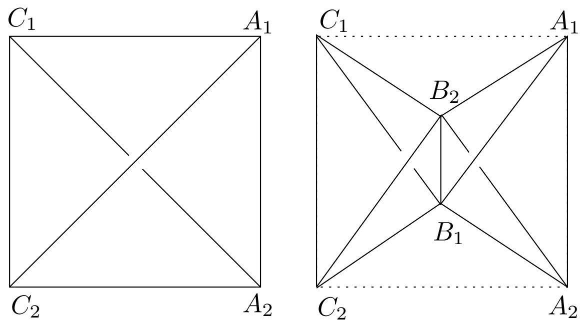

Figure 2.1

The 6 tetrahedra which partition Ω are glued together along common faces, in the following cyclic pattern.

$$\begin{tikzcd}
 & (A_1B_1C_2B_2) \arrow[dr] & \\
(C_1A_1C_2B_2) \arrow[ur] & & (A_1B_1A_2C_2) \arrow[d] \\
(C_1A_1B_2A_2) \arrow[u] & & (B_1C_1A_2C_2) \arrow[dl] \\
 & (B_1C_1B_2A_2) \arrow[ul] &
\end{tikzcd}$$

The action of σ translates this cycle of tetrahedra one third of the way around. A study of this pattern led us to the general case.

3 Theorem 1.1

3.1 The modular tiling

We use the disk model of $\boldsymbol{H}^2$. By slight abuse of terminology, we still say that $PSL_2(\boldsymbol{Z})$ acts on this model. Technically $PSL_2(\boldsymbol{Z})$ acts on the upper half plane model and a conjugate of $PSL_2(\boldsymbol{Z})$ acts on the disk model.

$\boldsymbol{H}^2$ has a canonical (and familiar) tiling T by ideal triangles which is invariant under the action of $PSL_2(\boldsymbol{Z})$. We define T by saying that it is the orbit of an ideal triangle under the group generated by reflections in its own sides. See [5, page 298] for a beautiful picture. We call T the *modular tiling*. Let VT, ET, and FT respectively denote the set of ideal vertices, geodesic edges, and ideal triangles of T. We say that two elements of ET are *touching* if they are identical or share a common endpoint. We say that two elements of FT are *touching* if they are identical or share a common edge.

T defines an exhaustion of $\boldsymbol{H}^2$ by ideal polygons. Let $t_+ \in FT$ be some distinguished ideal triangle. Let $T_0 = \{t_+\}$ and inductively define T_{m+1} to be those ideal triangles of T which are touching ideal triangles of T_m. Then T_m is an ideal polygon with 3×2^m sides. The *combinatorial distance* between $\tau_1, \tau_2 \in FT$ is the number of edges in ET crossed by the geodesic segment which connects the centers of τ_1 and τ_2. The triangles in $T_{m+1} - T_m$ are those which are have combinatorial distance m from t_+.

Each $e \in ET$ bounds a unique open halfspace h_e which is disjoint from the interior of t_+. Given $x \in S^1$ we write $e|x$ if x is an accumulation point of h_e. The set of x such that $e|x$ is one of the two closed arcs on S^1 determined by the endpoints of e. Each $x \in S^1 - VT$ defines a unique maximal sequence $\{e_m\}$ of edges such that $e_m|x$ for all x and $h_{m+1} \subset h_m$ for all m. We call $\{e_m\}$ the *nesting sequence* for x.

Using the disk model of $\boldsymbol{H}^2$ we can put a metric on $\boldsymbol{H}^2 \cup S^1$ which makes it isometric to a closed Euclidean disk. The next result refers to this metric.

Lemma 3.1 *For any $\eta > 0$ there is a $\delta > 0$ such that: If $x_1, x_2 \in S^1$ are less than δ apart then there are touching edges $e_1, e_2 \in ET$ such that $e_1|x_1$ and $e_2|x_2$.*

Proof There is some m such that all edges of ∂T_m, which is an ideal polygon, have diameter less than η. We take δ to be the minimim distance on S^1 between vertices of this ideal polygon. □

3.2 Modular pattern basics

Throughout this chapter Γ will denote a modular pattern of geodesics.

Lemma 3.2 *Let Γ be a modular pattern of geodesics.*

(1) *The endpoint of a geodesic of* Γ *belongs to* $S^1 - VT$.

(2) *Two geodesics of* Γ *cannot share an endpoint.*

(3) *Each* $e \in ET$ *intersects only finitely many geodesics of* Γ.

Proof Let G be a finite index torsion-free subgroup of $PSL_2(\boldsymbol{Z})$. Then T/G is a tiling of the finite area surface $\Sigma = \boldsymbol{H}^2/G$ by ideal triangles. Each individual edge of T maps injectively onto an edge of T/G. The quotient Γ/G is a finite union of closed geodesics on Σ. To prove Item 1, suppose a geodesic $\gamma \in \Gamma$ has an endpoint $v \in VT$. Then γ/G exits every compact subset of Σ as it approaches the cusp point on Σ corresponding to v. Closed geodesics have finite length and hence don't do this. Item 2 follows from the general fact, applied to γ_1/G and γ_2/G, that two closed geodesics on a complete hyperbolic surface cannot have lifts which share exactly one endpoint. To prove Item 3, note that the set $(e/G) \cap (\Gamma/G)$ is finite by compactness. Since the map $e \to e/G$ is injective the set $e \cap \Gamma$ is also finite. □

Item 2 above shows that the relation $\sim$ defined in section 1 is an equivalence relation: The transitivity condition is vacuously satisfied.

Corollary 3.3 *There is some* m *such that: If* $\tau_1, \tau_2 \in FT$ *have combinatorial distance at least* m *then at most one geodesic of* Γ *intersects both* τ_1 *and* τ_2.

Proof By symmetry we can take $\tau_1 = t_+$, the distinguished triangle. If $\{t_m\}$ was a sequence of counterexamples to this lemma then there would be two geodesics of Γ intersecting both t_+ and t_m for all m. Taking a subsequence we can assume that t_m converges to some $x \in S^1$. There are only finitely many geodesics which intersect t_+. Hence, taking another subsequence, we get the same two geodesics intersecting t_m for all m. But then x would be an endpoint to both geodesics, contradicting Item 2 of Lemma 3.2. □

3.3 Simplicial correspondences

We continue with the notation established above. For each $e \in ET$ let Γ_e denote the set with the following description. An object is an element of Γ_e iff it is either an endpoint of e or a geodesic of Γ which crosses e. The cardinality of Γ_e is $|\Gamma|+1$, where $|\Gamma|$ is as in Theorem 1.1. This quantity is finite by Lemma 3.2 and independent of e by symmetry. As in the statement of Theorem 1.1 we let $n = |\Gamma|$.

Recall from section 2.1 that $S^n = \Delta_+ \cup \Delta_-$. We equip S^n with the piecewise Euclidean metric inherited from Δ_+ and Δ_-. As in section 1.1, a *simplex* of S^n is defined to be a sub-simplex of either Δ_+ or Δ_-. If Δ is an n–simplex of S^n there is a bijective map from Γ_e to $V\Delta$, the vertex set of Δ, because the two sets have the same cardinality. Compare our table at the end of section 1.

To each $e \in ET$ we assign a pair

$$\Psi(e) = (\Delta_e, \phi_e), \tag{10}$$

where Δ_e is an n–dimensional simplex of S^n and $\phi_e\colon \Gamma_e \to V\Delta_e$ is a bijection. When the map ϕ_e is not immediately under discussion we will sometimes abuse notation and write $\Delta_e = \Psi(e)$.

We say that Ψ is a *simplicial correspondence* for Γ if it satisfies the following 3 properties:

Property 1 For any $\epsilon > 0$ there are only finitely many simplices in the image of Ψ which have diameter greater than ϵ.

Property 2 Given e_1 and e_2 in ET we let $(\Delta_j, \phi_j) = \Psi(e_j)$. Suppose v_j is a vertex of Δ_j for $j = 1, 2$. Then $v_1 = v_2$ iff $\phi_1^{-1}(v_1) = \phi_2^{-1}(v_2)$. So, each vertex in the grand union $\Psi(ET)$ is labelled by a unique element of $\Gamma \cap VT$. Compare the table at the end of section 1.

Property 3 Let e_1, e_2, Δ_1 and Δ_2 be as in Property 2. Let $h_j = h_{e_j}$ for $j = 1, 2$, as defined in section 3.1. We require that $\mathrm{int}(\Delta_1) \subset \mathrm{int}(\Delta_2)$ iff $h_1 \subset h_1$ and $\mathrm{int}(\Delta_1) \cap \mathrm{int}(\Delta_2) = \emptyset$ iff $h_1 \cap h_2 = \emptyset$. We also require that $\partial\Delta_1 \cap \partial\Delta_2$ is the convex hull of their common vertices. So, the simplices have the same nesting properties as the open half spaces.

We will construct Ψ in section 3.6. First we want to explore the consequences of its existence.

3.4 Embedding the quotient

In this section we use Ψ to define an embedding $i\colon Q_\Gamma \to S^n$. Let $x \in S^1$. Referring to the notation of section 3.1, there is a sequence $\{e_m\}_{m=1}^\infty \in ET$ such that $e_m | x$ for all m. (This is true even if $x \in VT$, but there is not a unique maximal sequence in this case.) Let $\Delta_m = \Psi(e_m)$ and

$$\Psi_\infty(x) = \bigcap_{n=1}^\infty \Delta_m. \tag{11}$$

Lemma 3.4 Ψ_∞ *is well defined.*

Proof If $x \in VT$ then x is an endpoint of e_m for all m. Hence $\phi_m(x) \in V\Delta_m$ for all m. Hence $\Psi_\infty(x) = \phi_m(x)$, independent of m and the choice of $\{e_m\}$. If $x \in S^1 - VT$ then any sequence used to define $\Psi_\infty(x)$ is contained in the nesting sequence for x. The intersection in Equation 11 is nested, by Property 3, and is a single point, by Property 1. $\square$

Lemma 3.5 Ψ_∞ *identifies points on* S^1 *if and only if they are equivalent.*

Proof If $x, x' \in S^1$ are endpoints of a geodesic γ in Γ then by Lemma 3.2 we have $x, x' \in S^1 - VT$. Let $\{e_m\}$ and $\{e'_m\}$ be the nesting sequences for x and x' respectively. Then γ crosses e_m and e'_m for all m and Δ_m and Δ'_m share a vertex for all m, by Property 2. Hence $\Psi_\infty(x) = \Psi_\infty(x')$.

If $x, x' \in S^1$ are inequivalent then there are edges $e, e' \in ET$ such that

(1) $h_e \cap h_{e'} = \emptyset$.

(2) e and e' have no vertices in common.

(3) $e|x$ and $e|x'$.

(4) No geodesic of Γ crosses both e and e'.

If this was false then we could take a limit of a sequence of counterexamples and produce a geodesic of Γ whose endpoints were x and x'.

Now $\Gamma_e \cap \Gamma_{e'} = \emptyset$ by Items 2 and 4. By Property 2, the simplices Δ_e and $\Delta_{e'}$ corresponding to e and e' have no vertices in common. Hence $\Delta_e \cap \Delta_{e'} = \emptyset$ by Property 3 and Item 1. From the definition of Ψ and Item 3 we have $\Psi_\infty(x) \in \Delta_e$ and $\Psi_\infty(x') \in \Delta_{e'}$. Hence $\Psi_\infty(x) \neq \Psi_\infty(x')$. $\square$

Lemma 3.6 Ψ_∞ *is continuous.*

Proof Let $\|\cdot\|$ denote the diameter in the piecewise Euclidean metric on S^n and also the Euclidean diameter on $\boldsymbol{H}^2 \cup S^1$. Let $\epsilon > 0$ be given. By Property 1 there is some $\eta > 0$ such that: If $e \in ET$ satisfies $\|e\| < \eta$ then $\|\Delta_e\| < \epsilon/2$. Here $\Delta_e = \Psi(e)$. Let δ be as in Lemma 3.1. If $\mathrm{dist}(x_1, x_2) < \delta$ then there are touching $e_1, e_2 \in ET$ such that $e_j|x_j$ and $\|e_j\| < \eta$ for $j = 1, 2$. But then $\Psi_\infty(x_1)$ and $\Psi_\infty(x_2)$ are contained in simplices Δ_1 and Δ_2 which by Property 2 share at least one vertex. Moreover $\|\Delta_j\| < \epsilon/2$. Hence $\mathrm{dist}(\Psi_\infty(x_1) - \Psi_\infty(x_2)) < \epsilon$. $\square$

Define

$$\Lambda = \Psi_\infty(S^1). \tag{12}$$

Combining the last two results we see that Ψ_∞ factors through a continuous bijection $i\colon Q_\Gamma \to \Lambda$. A continuous bijection from a compact space to a Hausdorff space is a homeomorphism. Thus i is a homeomorphism. This is our embedding from Theorem 1.1.

Here we give a useful characterization of Λ.

Lemma 3.7

$$\Lambda = \bigcap_{m=0}^{\infty} \Lambda_m \qquad \text{where} \qquad \Lambda_m = \bigcup_{e \in \partial T_m} \Delta_e. \tag{13}$$

Proof We have $\Lambda \subset \Lambda_m$ by Property 3 and the definition of Ψ. Any $y \in \bigcap \Lambda_m$ is contained in an infinite nested sequence $\{\Delta_m\}$ of simplices. By Property 3 the corresponding sequence $\{e_m\}$ is such that $e_m|x$ for some $x \in S^1$ and for all m. Thus $\Psi_\infty(x) = y$. Hence $\bigcap \Lambda_m \subset \Lambda$. □

Remark As above we set $\Psi(e) = (\Delta_e, \phi_e)$. By Property 2 all the local maps $\{\phi_e|\ e \in ET\}$ piece together to give a global bijection:

$$\left[\Gamma \cup VT = \bigcup_{e \in ET} \Gamma_e\right] \qquad \overset{\phi}{\longleftrightarrow} \qquad \left[V\Psi = \bigcup_{e \in ET} V\Delta_e\right]$$

If $x \in VT$ then $\Psi_\infty(x) = \phi(x)$. If x is an endpoint of a geodesic γ of Γ then $\Psi_\infty(x) = \phi(\gamma)$. Therefore Λ is the closure of $V\Psi$.

3.5 Block networks

Let Ω_+ be the modular block from section 2. We will only use modular blocks in S^n which have the following definition: Let $\Delta \subset S^n$ be a simplex. Let $A\colon \Delta_+ \to \Delta$ be an affine isomorphism. Our new modular block is $A(\Omega_+)$. The outer terminal is $\partial\Delta$. Every two modular blocks in S^n (that we use) are affinely equivalent. A maps the canonical triangulation of Ω_+ to a canonical triangulation of $A(\Omega_+)$. Given two modular blocks $\Omega_1, \Omega_2 \subset S^n$ we let $\mathrm{Map}(\Omega_1, \Omega_2)$ be the set of triangulation-respecting PL maps from Ω_1 to Ω_2.

Given $\tau \in FT$ we define

$$\Gamma_\tau = \bigcup_{e \in \partial\tau} \Gamma_e. \tag{14}$$

For each edge $e \in ET$, the set Γ_e has $2k$ elements. Each Γ_e shares k of its elements with another $\Gamma_{e'}$. Hence Γ_τ has $3k$ elements. An n–dimensional modular block Ω also has $3k$ vertices. Let $g \in PSL_2(\boldsymbol{Z})$ be the element which cycles the 3 edges of τ in counterclockwise order. Let σ be the 3–fold PL symmetry of Ω. We say that a *τ–labelling* of Ω is a bijection $\phi\colon \Gamma_\tau \to V\Omega$ which satisfies $\phi \circ g = \sigma \circ \phi$. Here $V\Omega$ is the vertex set of Ω.

Lemma 3.8 *If Ω_1 and Ω_2 are two τ–labelled modular blocks then there is a unique element of* $\mathrm{Map}(\Omega_1, \Omega_2)$ *which carries the one labelling to the other. This element is affine if it matches up the outer terminals.*

Proof Composing with affine maps we reduce to the case $\Omega_1 = \Omega_2 = \Omega_+$. Note first that $\mathrm{Map}(\Omega_+, \Omega_+)$ is quite large: Any permutation π of the the k–element set $A = \{A_j\}_{j=1}^k \subset V\Omega_+$ extends to an element $I_\pi \in \mathrm{Map}(\Omega_+, \Omega_+)$ which is an isometry. The map I commutes with σ, the 3–fold symmetry of Ω_+, and (hence) permutes the indices of the vectors in B and C in the same way it permutes the indices of the vectors in A.

Let ϕ and ϕ' be two τ–labellings of Ω_+. Let e_0, e_1, e_2 be the three edges of τ and let $\Delta_0, \Delta_1, \Delta_2$ be the simplices associated to Ω_+. Let $\Gamma_j = \Gamma_{e_j}$. Let $S = \Gamma_0 \cap \Gamma_1$. Composing ϕ with σ^a for some $a \in \{0, 1, 2\}$ we can assume that $\phi(\Gamma_j) = V\Delta_j = \phi'(\Gamma_j)$ for $j = 0, 1, 2$. By symmetry ϕ and ϕ' are determined by their action on S. Also $\phi(S) = \phi'(S)$. Hence there is some permutation π of $A = V\Delta_0 \cap V\Delta_1$ such that $\phi|_S = \pi \circ \phi'|_S$. Referring to the extension of π discussed above, we have $\phi = I_\pi \circ \phi'$. □

We say that a *block network* is an assignment $\tau \to \Omega[\tau]$, for each $\tau \in FT$. Here $\Omega[\tau]$ is a τ–labelled modular block. We require that $\Omega[t_+] = \Omega_+$ and

(1) $\Omega[\tau_1]$ and $\Omega[\tau_2]$ have disjoint interiors for all $\tau_1 \neq \tau_2$.

(2) $\Omega[\tau_1]$ and $\Omega[\tau_2]$ share a common terminal if τ_1 and τ_2 share an edge.

(3) $\Omega[\tau_1]$ and $\Omega[\tau_2]$ share a common vertex v if and only if the τ_1 label of v coincides with the τ_2 label of v.

Lemma 3.9 *There exists a block network for Γ.*

Proof We choose an enumeration $t_+ = t_0, t_1, t_2, ...$ of the ideal triangles of FT with the following property: For any $w \geq 1$, each t_w shares an edge e with some t_v for some $v < w$. We will define $\Omega_w = \Omega[t_w]$ inductively. We define $\Omega_0 = \Omega_+$ as we must. We choose some t_0–labelling for Ω_0. Note that $S^n - \Omega_0$

consists of 3 disjoint open simplices: $\mathrm{int}(\Delta_-)$ and $\mathrm{int}(\Delta_1)$ and $\mathrm{int}(\Delta_2)$. We call these simplices *holes*. Each edge of t_0 corresponds to a hole.

Suppose that $\Omega_0, ..., \Omega_{w-1}$ have been defined, and each edge of the polygon $P_w = \partial(t_0 \cup ... \cup t_{w-1})$ is associated to an open simplex—i.e. a hole—of $S^n - \bigcup_{j=1}^{w-1} \Omega_j$, There is some edge e of P_w which bounds t_w and some $v < w$ such that t_v and t_w share e as an edge. Let Δ be the simplex which is the closure of the corresponding hole in S^n. Note that $V\Delta$ is already labelled by elements of Γ_e. The labelling comes from the t_v–labelling of Ω_v, which has $\partial\Delta$ a terminal. First we choose a t_w–labelling of Ω_0 such that the outer terminal $\partial\Delta_+$ is labelled by elements of Γ_e. Next we choose the unique affine isomorphism $A\colon \Delta_+ \to \Delta$ which matches the t_w–labelling of Δ_+ with the t_v–labelling of Δ. We define $\Omega_w = A(\Omega_0)$. We use A to give Ω_w a t_w–labelling. The hole Δ has been plugged up but the two inner terminals of Ω_w bound two new holes.

Our construction only identifies vertices when they correspond to the same object of $\Gamma \cup VT$. No vertices are identified by accident because of the way the blocks are nested. These same nesting properties show that all the blocks have disjoint interiors. Thus we have constructed a block network. $\square$

Remark The axioms for block networks imply that any block network for Γ can be constructed by our inductive process. Once we determine the t_0–labelling the rest of the construction is forced. Different t_0–labellings produce geometrically identical networks, but with the labels permuted.

3.6 Putting it together

Let $\Omega[*]$ be our block network. We define $\Psi(e) = (\Delta_e, \phi_e)$, where Δ_e is the relevant simplex of $\Omega[\tau]$ and ϕ_e is the restriction of the τlabelling to Γ_e. Here τ is one of the two triangles which has e as an edge. From the block network axioms, either choice of τ gives the same map.

Lemma 3.10 *Ψ is a simplicial correspondence.*

Proof Properties 2 and 3 are immediate from our construction. It suffices to check property 1. Given two edges $e, e' \in ET$ we write $e \to_1 e'$ if $h_e \subset h_{e'}$ and if e and e' bound a common ideal triangle of T. We inductively define $e \to_{(m+1)} e'$ iff $e \to_m e''$ and $e'' \to_1 e'$. We let $\Delta_e = \Psi(e)$ and $\Delta_{e'} = \Psi(e')$. By Corollary 3.3, Property 2, and Property 3, there is some m such that: If $e \to_m e'$ then $\partial\Delta_e \cap \partial\Delta_{e'}$ is at most a single point. We fix m.

Let $\mathcal{S}$ denote the set of pairs of simplices of the form $(\Delta_e, \Delta_{e'})$, where $e \to_m e'$. We say that two pairs (Δ_1, Δ_1') and (Δ_2, Δ_2') in $\mathcal{S}$ are *equivalent* if there is an affine map which carries one pair to the other. Modulo $PSL_2(\boldsymbol{Z})$ there are only finitely many pairs (e, e') with $e \to_m e'$. Thus, by the affine naturality in our construction, $\mathcal{S}$ contains finitely many equivalence classes.

Let $\Delta = \Delta_+$, the standard simplex. For each equivalence class in $\mathcal{S}$ we define a *model pair* (Δ, Δ'), affine equivalent to any member of the equivalence class. If Property 1 fails 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. Such a nested sequence exists by Property 3. 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 infinitely often.

Let λ_j' be the longest edge of Δ_{j+1}. Let λ_j be the longer line segment obtained by intersecting Δ_j with the line containing λ_j'. Since $\bigcap \Delta_j$ does not shrink to a point, $\text{length}(\lambda_j') \not\to 0$. Since Δ_j and Δ_{j+1} converge to each other as $j \to \infty$, we have $\text{length}(\lambda_j') \to \text{length}(\lambda_j)$. 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 $\text{length}(\delta_j') \to \text{length}(\delta_j)$ and δ_j' is an edge of Δ' for all j. Everything takes place on the same model so the set of possible pairs (δ_j', δ_j) is finite. Hence $\delta_j = \delta_j'$ for large j. Hence $\partial\Delta$ and $\partial\Delta'$ have two distinct points in common, contradicting the choice of m. □

The work in section 3.4 gives us our embedding. Now we construct the representation from Theorem 1.1. Let $g \in PSL_2(\boldsymbol{Z})$. Let τ be an ideal triangle of T. Let $\tau' = g(\tau)$. Let $\phi\colon \Gamma_\tau \to \Omega[\tau]$ be the τ–labelling of $\Omega[\tau]$. Let $\phi'\colon \Gamma_{\tau'} \to \Omega[\tau']$ be the τ'–labelling of $\Omega[\tau']$. By Lemma 3.8 there is a unique $\rho(g,\tau) \in \text{Map}(\Omega[\tau], \Omega[\tau'])$ such that $\rho(g,\tau) \circ \phi = \phi'$. In other words $\rho(g,\tau)$ maps the vertex of $\Omega[\tau]$ labelled by the object x to the vertex of $\Omega[\tau']$ labelled by the object $x' = g(x)$. This *intertwining property* implies that $\rho(g,\tau_1)$ and $\rho(g,\tau_2)$ agree on any common vertices. Since $\Omega[\tau_1] \cap \Omega[\tau_2]$ is contained in a single simplex, on which both our maps are affine, we see that $\rho(g,\tau_1) = \rho(g,\tau_2)$ on $\Omega[\tau_1] \cap \Omega[\tau_2]$. Letting $\widehat{\Omega} = \bigcup \Omega[\tau]$ we see that the $\rho(g,\tau)$ maps piece together to give a continuous map $\rho(g)\colon \widehat{\Omega} \to \widehat{\Omega}$. The intertwining property gives

$$\rho(g_1 g_2) = \rho(g_1)\rho(g_2). \tag{15}$$

We now show that $\rho(g)$ extends to an element of $\text{PL}(S^n)$. There is some m such that $t_+ \subset g(T_m)$. If e is an edge of $\tau_1 \in FT - T_m$ then $h_e \cap T_m = \emptyset$. Therefore $\tau_1 \notin g(h_e)$. Therefore $g(h_e) = h_{g(e)}$. Therefore $\rho(g,\tau_1)$ identifies the outer terminals of the two blocks and by Lemma 3.8 is affine. If τ_2 is an ideal triangle touching τ_1 and contained in h_e then $\Omega[\tau_1]$ and $\Omega[\tau_2]$ intersect along the outer terminal $\partial\Delta$ of $\Omega[\tau_1]$. Since two affine maps are determined by their action on a simplex we see that $\rho(g,\tau_1)$ and $\rho(g,\tau_2)$ are restrictions of the same affine map. Repeating this argument with τ_2 replacing τ_1, etc., we see inductively that $\rho(g)$ is affine on all blocks contained in Δ. We extend $\rho(g)$ by making it affine on all of Δ. Since there are only finitely many edges of T_{m+1} we see that the extension of $\rho(g)$ is PL on the set Λ_{m+1} defined in Lemma 3.7. There are only finitely many blocks not contained in Λ_{m+1} and $\rho(g)$ is PL on each one. In summary, $\rho(g)$ is a PL map. Equation 15 shows that $\rho(g)$ has the inverse $\rho(g^{-1})$ which is also PL. Hence $\rho(g) \in \mathrm{PL}(S^n)$.

Equation 15 says that the map $g \to \rho(g)$ is a homomorphism. Every $\rho(g)$ acts nontrivially on some block. Hence ρ is a monomorphism. Λ is the closure of the block network vertices. Hence $H = \rho(PSL_2(\boldsymbol{Z}))$ preserves Λ. From the remark at the end of section 3.4, the map Ψ conjugates the minimal action of $PSL_2(\boldsymbol{Z})$ on $\partial\boldsymbol{H}^2$ to a minimal action of H on Λ. The triangulations of all the blocks piece together to give a partition of $S^n - \Lambda$ by punctured simplices. Corollary 3.2 and the local finiteness of the modular tiling imply that our partition by punctured simplices is locally finite. H permutes this partition and hence acts properly discontinuously on $S^n - \Lambda$. In short Λ is the limit set of H. Our proof of Theorem 1.1 is done.

4 Modified blocks

4.1 Partial prisms

Let $n = 2k-1$ as in previous chapters. A *convex cone* in $\boldsymbol{R}^n$ is a closed convex subset $C \subset \boldsymbol{R}^n$, contained in a halfspace, which is closed under taking non-negative linear combinations. C is *generated* by the set Σ if $C = \{\lambda\Sigma \mid \lambda \geq 0\}$. We call C a *simplex-cone* if C is generated by an $(n-1)$–simplex which does not contain 0. We also insist that C is n–dimensional.

Let C be a simplex-cone. Let $H \subset \boldsymbol{R}^n$ be a codimension 1 hyperplane which does not contain 0. We say that H *cuts* C if $H \cap C$ is an $(n-1)$–simplex. In this case C is generated by $H \cap C$. If H cuts C we set $\Sigma = C \cap H$ and let $[C,H] = \{\lambda\Sigma | \ \lambda \leq 1\}$. With this definition, $[C,H]$ is an n–simplex, one of whose vertices is 0, and whose other vertices are the vertices of Σ. We say that a *partial prism* is a set isometric to a set of the form

$$\Pi = \mathrm{closure}([C,H_1] - [C,H_0]). \tag{16}$$

where H_0 and H_1 cut C and $[C, H_0] \subset [C, H_1]$. We call $C \cap H_0$ the *inner boundary* of Π and we call $C \cap H_1$ the *outer boundary* of Π. Note—and this is crucial for our constructions—that the inner and outer boundaries can share vertices in common or even coincide. In all cases there is a canonical bijection between the inner boundary vertices and the outer boundary vertices: The matched vertices lie on the same line through 0.

$\partial\Pi$ consists of two $(n-1)$–simplices—the inner and outer boundaries—and some $(n-1)$–dimensional partial prisms. This lets us define a canonical PL involution of Π which interchanges the inner and outer boundaries. If $n = 1$ then Π is an interval or a point, and our involution reverses the interval or fixes the point depending on the case. In general the PL involution is the cone, to the center of mass of Π, of the PL involution which is defined on each partial prism of $\partial\Pi$ and which swaps inner and outer boundary components. We call this map the *canonical involution*.

We also can define a canonical triangulation of Π. If Π is a simplex then we use Π itself as the triangulation. Otherwise we triangulate $\partial\Pi$ (by induction) and then cone the resulting triangulation to the center of mass of Π. We call this the *canonical triangulation* of Π. The canonical PL involution is affine when restricted to each simplex in the canonical triangulation. The important point about the canonical triangulation is that it has this 2–fold symmetry.

4.2 Separators

Let $n = 2k - 1$ as above. Let $\langle \cdot \rangle$ be the convex hull operation. We say that a *weighted simplex* is an n–simplex Δ together with a map $S\colon V\Delta \to (0, 1]$. Let $v_1, ..., v_{n+1}$ be the vertices of Δ. Let $S_i = S(v_i)$. Let

$$v_i^* = S_i v_i + (1 - S_i)\beta_S; \qquad \beta_S = \frac{\sum_{i=1}^{n+1} S_i v_i}{\sum_{i=1}^{n+1} S_i} \in \Delta \tag{17}$$

Note that $v_i^* = v_i$ if and only if $S_i = 1$. In all cases v_i^* is contained in the half-open interval $(\beta_S, v_i]$ which joins β_S to v_i. Hence $v_1^*, ..., v_{n+1}^*$ are in general position. Finally, we define

$$\Delta_S = \left\langle \bigcup_{i=1}^{n+1} v_i^* \right\rangle; \qquad [\Delta, S] = \mathrm{closure}(\Delta - \Delta_S) \tag{18}$$

We call $[\Delta, S]$ a *separator*. We call $\partial\Delta$ and $\partial\Delta_S$ respectively the *inner* and *outer boundaries* of $[\Delta, S]$. For each codimension 1 face Σ_S of Δ_S there is a unique codimension 1 face Σ of Δ such that $\Sigma_S \subset \langle \beta_S \cup \Sigma \rangle$. The set

$$\langle \beta_S \cup \Sigma \rangle - \langle \beta_S \cup \Sigma_S \rangle \tag{19}$$

is a partial prism. Therefore $[\Delta, S]$ is canonically partitioned into $n+1$ partial prisms.

The canonical involutions on the individual prisms piece together to produce a canonical involution of $[\Delta, S]$ which swaps $\partial\Delta$ and $\partial\Delta_S$. This involution is affine on each of $\partial\Delta$ and $\partial\Delta_S$. The canonical triangulations of the individual prisms piece together to give a *canonical triangulation* of the separator. The canonical involution of the separator is affine when restricted to each simplex in the canonical triangulation. The combinatorial structure of the canonical triangulation depends entirely on the set $S^{-1}(1) \subset V\Delta$. Any permutation of $V\Delta$ which respects this set extends to a PL automorphism of $[\Delta, S]$ which is affine on each simplex in the triangulation and also affine on the boundaries. This observation underlies Lemma 5.1 below.

If $\psi\colon \Delta \to \Delta'$ is an affine isomorphism which carries S to S' then $\psi([\Delta, S]) = [\Delta', S']$. Thus the separator construction is affinely natural. Also $[\Delta, S]$ varies continuously with S. Finally, a short computation reveals that β_S is the barycenter of Δ_S.

4.3 Warped blocks

Say that a *weighted block* is a block Ω equipped with a weighting of its vertices, $S\colon V\Omega \to (0,1]$. Let (Ω, S) be a weighted block. Let $\partial\Delta_0$ be the outer terminal of Ω, so that $\Omega \subset \Delta_0$. Let $\partial\Delta_1$ and $\partial\Delta_2$ be the inner terminals of Ω. Let $v_1, ..., v_{n+1}$ be the vertices of Δ_0. Every point of Δ_0 has the form

$$x = \sum_{i=1}^{n+1} \lambda_i v_i; \qquad \text{where} \qquad \sum_{i=1}^{n+1} \lambda_i = 1. \tag{20}$$

Let $S_i = S(v_i)$. We define

$$P_S(x) = \frac{\sum_{i=1}^{n+1} S_i \lambda_i v_i}{\sum_{i=1}^{n+1} S_i \lambda_i} \in \Delta_0. \tag{21}$$

The map P_S is not a linear map. However it is a projective automorphism of Δ_0. In particular P_S permutes the set of simplices contained in Δ_0. We define

$$\Omega_S = P_S(\Omega). \tag{22}$$

The terminals of Ω_S are the three simplex boundaries

$$\partial\Delta_0; \qquad \Delta_{1,S} = P_S(\partial\Delta_1); \qquad \Delta_{2,S} = P_S(\partial\Delta_2). \tag{23}$$

P_S maps the triangulation of Ω to a combinatorially equivalent triangulation of Ω_S. Thus Ω_S has a canonical triangulation. There is a canonical PL homeomorphism $W_S\colon \Omega \to \Omega_S$ which is affine on each simplex of the triangulations. W_S conjugates the 3–fold PL symmetry σ of Ω to a 3–fold PL symmetry σ_S of Ω_S. By construction σ_S is affine when restricted to each of the terminals of Ω_S.

We call Ω_S a *warped block*. The map W_S sets up a canonical bijection between $V\Omega$ and $V\Omega_S$. In this way we transfer the map $S\colon V\Omega \to (0,1]$ to a map $S\colon V\Omega_S \to (0,1]$. In other words the vertices of Ω_S are naturally weighted. Note that S restricts to give a weighting to each of the terminals of Ω_S. For instance, we have the restriction map $S\colon V\Delta_{1,S} \to (0,1]$.

If (Ω_1, S_1) and (Ω_2, S_2) are weighted blocks and $T\colon \Omega_1 \to \Omega_2$ is an affine map such that $S_1 = S_2 \circ T$ then $T(\Omega_{S_1}) = \Omega_{S_2}$. This follows from the fact that T conjugates P_{S_1} to P_{S_2}, as can be seen from Equation 21. Our warping construction is affinely natural even though the map P_S is not itself affine.

4.4 Main constructions

General modified blocks Let (Ω, S) be as above. Let Δ_S be as in Equation 18. Let T_S be the affine map which carries Δ_S to Δ in such a way that $T_S(v_i^*) = v_i$ for all i. Note that $T_S([\Delta, S])$ is a separator whose inner boundary is Δ. We define

$$[\Omega, S] = \Omega_S \cup T_S([\Delta_0, S]) \cup [\Delta_{1,S}, S] \cup [\Delta_{2,S}, S]. \tag{24}$$

We have attached one separator to each terminal of Ω_S. We call $[\Omega, S]$ a *general modified block*. We call Ω_S the *core* of $[\Omega, S]$. Note that $[\Omega, S]$ again has three terminals; these are the free boundaries of the attached separators. The outer terminal is $T_S(\partial\Delta_0)$.

Remarks

(i) Note that $\partial\Delta_0$ is the inner boundary of $T_S([\Delta_0, S])$ whereas $\partial\Delta_j$ is the outer boundary of $[\Delta_j, S]$. From a PL standpoint this asymmetry in our construction disappears: Each separator has its canonical involution which turns it inside out.

(ii) $[\Omega, S]$ is not necessarily a subset of S^n. The problem is that the outer boundary of $T_S([\Delta_0, S])$ might be so large that it is not contained in one of the two unit simplices comprising S^n. This difficulty will be handled in section 5 in an automatic way. Our construction will only use modified blocks which are contained in S^n.

(iii) Given our definitions in Equations 18 and 24, each vertex $v \in V\Omega_S$ corresponds to two vertices $v_1, v_2 \in V[\Omega, S]$ and we have $v_1 = v_2$ if and only if $S(v) = 1$. The same remarks apply to $]\Omega_+, S[$ below.

Special modified blocks Suppose now that Ω_+ is the modular block constructed in section 2. Then $\Omega_+ = \text{closure}(S^n - \Delta_- - \Delta_1 - \Delta_2)$. This follows from the fact that the outer terminal of Ω_+ is $\partial\Delta_+$, and $S^n = \Delta_+ \cup \Delta_-$. The weighting S gives a map $S\colon V\Delta_- \to (0,1]$ as well as the maps $S\colon V\Delta_{j,S} \to (0,1]$ for $j = 1, 2$. We define

$$]\Omega_+, S[= (\Omega_+)_S \cup [\Delta_-, S] \cup [\Delta_{1,S}, S] \cup [\Delta_{2,S}, S]. \tag{25}$$

The first separator is contained in Δ_-. We call $]\Omega_+, S[$ a *special modified block*. We call $(\Omega_+)_S$ the *core* of $]\Omega_+, S[$. The free boundaries of the separators are the terminals.

4.5 Degeneration

Let $n_1 < n_2$ be two integers. Let $\{\Delta_m\}$ denote a sequence of n_2–simplices. Let Δ' be an n_1–simplex. We say that Δ_m *converges barycentrically* to Δ' if some collection of n_1 vertices of Δ_m converges to the vertices of Δ' as $m \to \infty$ and the remaining vertices of Δ_m converge to the barycenter of Δ'. (We shall always have a consistent labelling of the vertices.) Referring to section 4.2:

Lemma 4.1 *Let Δ be an n_2–simplex and let Δ' be an n_1–simplex face of Δ. Let $S_m\colon V\Delta \to (0,1]$ and $S'\colon V\Delta' \to (0,1]$ be such that $S_m(v) \to S'(v)$ if $v \in V\Delta'$ and $S_m(v) \to 0$ otherwise. Δ_{S_m} converges barycentrically to $\Delta'_{S'}$.*

Proof Equation 17 extends continuously to the case when some (but not all) of the S_i are zero. Extending S' by the 0–map we have $S' = \lim S_m$. When $S'_i = 0$ we have $v_i^* = \beta_{S'}$, the barycenter of $\Delta'_{S'}$. □

Let Ω_+^j denote the n_j–dimensional block from the Block Lemma. In general we use the notation X^j to refer to an object associated to Ω_+^j though sometimes we simplify the notation. Referring to Equation 2 there is a natural embedding $i\colon S^{n_1} \to S^{n_2}$ defined by $i(A_j^1) = A_j^2$ and $i(C_j^2) = C_j^2$ for $j = 1, ..., k_1$. Suppose $S^1\colon V\Omega_+^1 \to (0,1]$. We define $(\Omega', S') = i(\Omega_+^1, S_1)$ and $\Omega = \Omega_+^2$. Suppose

$S_m\colon V\Omega \to (0,1]$ is a sequence of maps such that $S_m(v) \to S'(v)$ if $v \in V\Omega'$ and $S_m(v) \to 0$ otherwise. Let $]\Omega, S_m[$ and $[\Omega, S_m]$ be the special and general modified blocks based on (Ω, S_m). Say that a *filled-in terminal* of a modified block is a simplex bounded by a terminal. The following result is the key to Theorem 1.4.

Lemma 4.2 *The filled-in terminals of* $]\Omega, S_m[$ *converge barycentrically to the filled-in terminals of* $]\Omega', S'[$.

Proof By Equations 18 and 25 the filled-in terminals of $]\Omega, S_m[$ are $(\Delta^2_-)_{S_m}$ and $(\Delta^2_{j,S_m})_{S_m}$. Let $\Delta' = i(\Delta^1_+)$ and $\Delta'_- = i(\Delta^1_-)$ and $\Delta'_j = i(\Delta^1_j)$. The filled-in terminals of $]\Omega', S'[$ are $(\Delta'_-)_{S'}$ and $(\Delta'_{j,S'})_{S'}$. In all cases, $j \in \{1,2\}$. Now, S_m, Δ^2_+, S' and Δ' are as in Lemma 4.1. Hence $(\Delta^2_-)_{S_m}$ converges barycentrically to $(\Delta'_-)_{S'}$. A direct calculation (which we did numerically on examples to be sure) shows that the first $2k_1$ vertices of $\Delta^2_{j,S_m} = P_{S_m}(\Delta^2_j)$ converge to the vertices of $\Delta'_{j,S'} = P_{S'}(\Delta'_j)$. Lemma 4.1 finishes the proof in this case. □

5 The rest of the results

5.1 Modified correspondences

Suppose that Γ is a modular pattern and $\Gamma' \subset \Gamma$ is a modular sub-pattern. We define Γ'_e just as we defined Γ_e. We have $\Gamma'_e \subset \Gamma_e$ for all $e \in ET$. To each $e \in ET$ we assign a pair $\Psi'(e) = (\Delta_e, \phi_e)$, where Δ_e is an n–dimensional simplex of S^n and $\phi_e\colon \Gamma_e \to V\Delta_e$ is a bijection. This is as in section 3.3. We say that Ψ' is a *modified simplicial correspondence* for the pair (Γ', Γ) if it satisfies the Properties 1 and 3 for simplicial correspondences and

Property 2′ Given e_1 and e_2 in ET we let $(\Delta_j, \phi_j) = \Psi'(e_j)$. Suppose v_j is a vertex of Δ_j for $j = 1,2$. Then $v_1 = v_2$ iff $\phi_1^{-1}(v_1) = \phi_2^{-1}(v_2)$ *and* the common object $\phi_j^{-1}(v_j)$ belongs to Γ'_e.

We define the map $\Psi'_\infty\colon S^1 \to S^n$ just as in Equation 11. Lemmas 3.4, 3.6 and 3.7 work exactly the same way for Ψ'_∞ as they do for Ψ_∞. Property 2′ causes a change in Lemma 3.5. The same argument in Lemma 3.5 proves that Ψ' identifies points on S^1 if and only if they are the common endpoints of a geodesics in Γ'. Thus Ψ'_∞ factors through an embedding of $Q_{\Gamma'}$ into S^n.

Remark The remark at the end of section 3.4 needs to be modified in the setting here. Property $2'$ gives a bijection between $\Gamma' \cup VT$ and a certain subset $V'\Psi' \subset V\Psi'$ of the block vertices. $\Lambda = \Psi'(S^1)$ is the closure of $V'\Psi'$.

5.2 Modified block networks

Each modified block has a canonical triangulation, obtained from the triangulations on the core and on the separators. Suppose that Ξ_1 and Ξ_2 are modified blocks with symmetries σ_1 and σ_2. Let $\mathrm{Map}(\Xi_1, \Xi_2)$ be the set of triangulation-respecting PL maps from Ξ_1 to Ξ_2. Suppose that Ξ_j has a weighted core $(\Omega_j)_{S_j}$ for $j = 1, 2$. We say that the bijection ψ: $V((\Omega_1)_{S_1}) \to V((\Omega_2)_{S_2})$ between the core vertex sets is a *perfect matching* if $\psi^{-1} \circ \sigma_2 \circ \psi = \sigma_1$ and $S_2 \circ \psi = S_1$. In other words, ψ is symmetry-respecting and weight-respecting.

Lemma 5.1 *A perfect matching ψ extends to an element of* $\mathrm{Map}(\Xi_1, \Xi_2)$. *When Ξ_1 and Ξ_2 are general modified blocks, this extension is affine if it matches up the outer terminals.*

Proof The combinatorial structure of the separators of Ξ only depends on S. The combinatorially identical triangulations on Ξ_1 and Ξ_2 define the extension of ψ. When the outer terminals are matched up, the extension of ψ to the cores is an affine map $\widehat{\psi}$. This follows from the affine naturality of the warping process. It follows from Equation 17 that the map $\widehat{\psi}$ maps the separators of Ξ_1 to the separators of Ξ_2. This follows from the affine naturality of the separator construction. $\square$

The rest of our constructions depend on some $f \in \square(\Gamma)$, which we fix throughout the discussion. We say that the elements of VT have weight 1. This convention, together with f, assigns weights to each element of Γ_τ, the set in Equation 14. Let $[\Omega, S]$ be a modified block. Let $\tau \in FT$ be an ideal triangle. We say that a *τ–labelling* of $[\Omega, S]$ is a bijection ϕ: $\Gamma_\tau \to V\Omega_S$ such that $\phi \circ g = \sigma \circ \phi$ *and* $S \circ \phi = f$. All maps above have Γ_τ as their range. As in section 3.5, the element g is the order 3 stabilizer of τ which cycles the edges counterclockwise. σ is the order 3 PL symmetry of $[\Omega, S]$. So, ϕ carries the weights of f to the weighting of Ω_S. We make the same definitions for $]\Omega_+, S[$.

We have labelled $V\Omega_S$, because $V[\Omega, S]$, the actual vertex set of $[\Omega, S]$, generally has more vertices than Γ_τ has elements. Here we describe an *induced labelling* of $V[\Omega, S]$. Let $\partial\Delta$ be one of the terminals of $[\Omega, S]$. Then one of the terminals $\partial\Delta'$ of Ω_S is such that $\partial\Delta$ and $\partial\Delta'$ form the boundary of a separator of $[\Omega, S]$. As with all separators, there is a canonical bijection $\beta\colon V\Delta \to V\Delta'$. One of the three edges e bounding τ is such that $\phi^{-1}(V\Delta) = \Gamma_e$. We label the vertex $\beta(v) \in V\Delta$ by the pair $(\phi^{-1}(v), e)$. In this way, each vertex of $V[\Omega, S]$ is labelled by a pair (γ, e), where $\gamma \in \Gamma_e$ and e is an edge of τ. Given Remark (iii) in section 4.4, and our construction here, the induced labelling has the property that v_1 is labelled by a pair (γ, e_1) and v_2 is labelled by a pair (γ, e_2). Here γ is the element of Γ_τ which labels v. We have $v_1 = v_2$ if and only if $f(\gamma) = 1$. We call this the *separation principle*.

We define *modified block networks* just as we defined block networks in section 3.5, using modified blocks in place of blocks. The one twist is that $\Omega[t_+]$ is a special modified block and all the other $\Omega[\tau]$ are general modified blocks.

Lemma 5.2 *There exists a modified block network for $f \in \Box(\Gamma)$.*

Proof The proof is essentially the same as the one given in Lemma 3.9. Let $t_+ = t_0, t_1, t_2, ...$ be as in Lemma 3.9. We need to construct modified blocks $\Omega_0, \Omega_1, \Omega_2 ...$, where $\Omega_j = \Omega[t_j]$. We set $\Omega_0 =]\Omega_+, S[$ as we must. At the induction step we choose the affine map A which takes the outer terminal of $[\Omega_+, S]$ to $\partial\Delta$, the terminal corresponding to the edge e of t_v, in such a way as to respect the labellings. □

Remark As in section 3.5 the modified block network is unique up to the choice of the t_0–labelling. However, if we base our construction on some general system of weights that is not invariant under $PSL_2(\boldsymbol{Z})$, as we do in the proof of Theorem 1.3 below, then there are potentially as many different geometric types of network as their are G equivalence classes of edges in ET. Here G is the symmetry group of f (which we will take to be a finite index subgroup of $PSL_2(\boldsymbol{Z})$).

5.3 Proof of Theorem 1.2

We continue with the notation from above. We set $\Gamma' = f^{-1}(1)$. We use our modified block network to define a modified correspondence for (Γ', Γ): For each edge $e \in ET$ we define $\Psi(e) = (\Delta_e, \phi_e)$, where Δ_e is the relevant boundary simplex of $\Omega[\tau]$ and ϕ_e is the labelling of Δ_e induced by the τ–labelling of $\Omega[\tau]$. Here τ is one of the two ideal triangles which has e in its boundary. Our construction guarantees that $\Psi'(e)$ is the same using either choice of τ. The nesting properties of the modified block network are the same as for the original block network. Hence Ψ' has property 3. The same argument as in section 3.6 shows that Ψ' has Property 1.

Property $2'$ follows from the Separation Principle. To see how this works, we consider our construction from a different point of view. We start with the block network for Γ. We then warp each block in the network. This changes the geometry of the network, but none of its combinatorial structure. The warped network still has property 2. Next, we split apart the terminals and insert separators—two per terminal because the terminal includes into two warped blocks. (We like to think of this as blowing air into the terminals.) Figure 5.1 shows a schematic picture.

The separators have the effect of splitting apart vertices which are labelled by geodesics γ in Γ which have weight less than 1. In the warped block network γ labels a single vertex. After the separators are added, there is an infinite list of vertices associated to γ. Each of these vertices has a label of the form (γ, e), where e is an edge of ET crossed by γ. If γ has weight 1, then all these infinitely many vertices coalesce into one. The separators do not affect these weight-1 vertices.

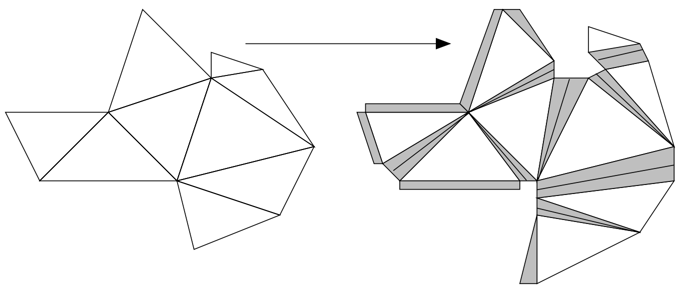

Figure 5.1

Thus our modified block network defines a modified simplicial correspondence Ψ' for (Γ', Γ). As in section 5.1 we have our embedding $i\colon Q_{\Gamma'} \to S^n$. We define $\Lambda_f = i(Q_{\Gamma'})$. The representation ρ_f is constructed exactly as in section 3, with Lemma 5.1 used in place of Lemma 3.8. The same argument as in section 3 shows that Λ_f is the limit set of ρ_f. The modified blocks and their symmetries are continuous functions of $f \in \square(\Gamma)$. Thus our two maps $\rho\colon \square(\Gamma) \to \mathrm{PL}(G, S^n)$ and $\Lambda\colon \square(\Gamma) \to [S^n]$ are continuous maps in the appropriate topologies.

5.4 Proof of Theorem 1.3

The modular group is hiding behind Theorem 1.3.

Lemma 5.3 *Any cusped finite volume hyperbolic surface is homeomorphic to a quotient of the form* $\boldsymbol{H}^2/G$*, where* G *is a finite index modular subgroup.*

Proof This is a well-known result. Every cusped surface has a triangulation into ideal triangles. Each edge of an ideal triangle has a center point, the fixed point set of the isometric involution of the triangle which stabilizes that edge. We can cut apart our surface and re-glue the ideal triangles so that the center points of the edges are matched. This changes the geometric structure but not the topology. The resulting surface then develops into the hyperbolic plane, onto the modular tiling. Thus the new surface, which is homeomorphic to the original, has the form $\boldsymbol{H}^2/G$ with $G \subset PSL_2(\boldsymbol{Z})$. $\square$

By Lemma 5.3 it suffices to consider the case of Theorem 1.3 where $\Sigma = \boldsymbol{H}^2/G$, so that $\pi_1(\Sigma) = G$, a finite index modular subgroup. Let Γ be the orbit of Γ' under $PSL_2(\boldsymbol{Z})$. Since G has finite index in $PSL_2(\boldsymbol{Z})$, we have that Γ is a modular pattern. We can define a modified simplicial correspondence for the pair (Γ', Γ) even when Γ' does not have complete modular symmetry. The definitions and results in section 5.1 go through word for word.

Let $f\colon \Gamma \to \{\frac{1}{2}, 1\}$ be the map defined by the rule $f(\gamma) = 1$ if $\gamma \in \Gamma'$ and $f(\gamma) = 1/2$ if $\gamma \in \Gamma - \Gamma'$. Even though f is not necessarily $PSL_2(\boldsymbol{Z})$–invariant we can define a modified block network for f. This network has G–symmetry rather than $PSL_2(\boldsymbol{Z})$–symmetry from the PL standpoint. The modified block network in turn defines a modified simplicial correspondence for (Γ', Γ). The same argument as in the proof of Theorem 1.2 (which is just adapted from section 3) gives a representation $\rho\colon G \to \mathrm{PL}(S^n)$ which has $i(Q_{\Gamma'})$ as its limit set.

5.5 Proof of Theorem 1.4

Given an affine map A let $\|A\| = \sup_v \|A(v)\|$ be the operator norm, with the sup being taken over unit vectors.

Lemma 5.4 *Let* Δ' *be an* n_1*–dimensional face of* Δ*, the unit* n_2*–simplex. Let* $\{A_m\}$ *be a sequence of affine maps of* $\boldsymbol{R}^{n_2}$*, with uniformly bounded operator norm, such that* $A_m|_{\Delta'}$ *converges to an affine injection* $A'\colon \Delta' \to \boldsymbol{R}^{n_2}$*. Let* $\{\Xi_m\}$ *be a sequence of* n*–simplices which converge barycentrically to some* n'*–simplex* $\Xi' \subset \Delta'$*. Then* $A_m(\Xi_m)$ *converges barycentrically to* $A'(\Xi')$*.*

Proof The first n_1 vertices of Ξ_m converge to the vertices of Ξ'. The bound on the operator norms guarantees that the first n_1 vertices of $A_m(\Xi_m)$ converge to $A'(\Xi')$. The remaining vertices of Ξ_m converge to the barycenter of Ξ'. Again, the bound on the operator norms guarantees the images of these remaining vertices under A_m converge to the barycenter of $A'(\Xi')$. □

We continue the notation from section 4.6 and also use the notation from Theorem 1.4. Let $]\Omega, S_m[$ and $[\Omega, S_m]$ be the special and general modified blocks based on $\Omega = {}_2\Omega_+$, corresponding to f_m. Thus $]\Omega, S_m[$ is the zeroth modified block in the modified block network for f_m and $[\Omega, S_m]$ is the general modified block used in the induction step of Lemma 5.2. We let S be the weighting on ${}_1\Omega_+$ that correponds to $f \in \square(\Gamma_1)$.

For each m we need to choose a t_+–labelling of $]\Omega, S_m[$. We pick the labelling so that the vertices in $A' \cup C'$ are labelled by objects associated to Γ_1. We can make the labellings independent of m, since only the weights vary with m. We can choose a t_+–labelling of $]_1\Omega, S[$ which is consistent with our t_+–labellings of $]\Omega, S_m[$. All the same remarks apply to $[\Omega, S_m]$ and $[_1\Omega, S]$. This sets things up so that $]\Omega, S_m[$ and $[\Omega, S_m]$ are as in Lemma 4.2.

For each m we have a modified block network $N_m \subset S^{n_2}$. We also have a modified block network $N \subset S^{n_1}$. Let $N' = i(N)$. The terminals of N_m are canonically bijective with the terminals of N'. Both are indexed by ET.

Lemma 5.5 *Each filled-in terminal of N_m converges barycentrically to the corresponding terminal of N' as $m \to \infty$.*

Proof Let $t_0, t_1, t_2...$ be as in Lemma 5.2. For ease of notation we suppress the dependence on m. Let Ω_j be the modified block associated to t_j when the construction is based on $f_m \in \square(\Gamma_2)$. Let ${}_1\Omega_j$ be the modified block associated to t_j when the construction is based on $f \in \square(\Gamma)$. Let $\Omega'_j = i({}_1\Omega_j)$. It suffices to show that the filled-in terminals of Ω_j converge barycentrically to the filled-in terminals of Ω'_j. For $j = 0$ this is exactly Lemma 4.2.

Suppose the result is true for $j = 1, ..., w-1$. We consider the case $j = w$. We adopt the notation from Lemma 3.9 and 5.2. Thus D is the outer filled-in terminal of $[\Omega_+, S_m]$ and A: $D \to \Delta$ is such that the two filled-in inner terminals of Ω_w are $A(\Delta_1)$ and $A(\Delta_2)$. Here Δ_1 and Δ_2 are the two inner filled-in terminals of $[\Omega_+, S_m]$. By induction Δ converges to barycentrically to one of the inner terminals of Ω'_v. Thus the outer filled-in terminal of Ω_w converges barycentrically to the outer filled-in terminal of Ω'_w. We just have to show that $A(\Delta_1)$ and $A(\Delta_2)$ converge barycentrically to the inner filled-in terminals of Ω'_w.

Note that $\Delta_+ \subset D$ and either $A(D) \subset \Delta_+$ or $A(D) \subset \Delta_-$. In either case A maps the standard unit simplex inside an isometric copy of itself. This bounds $\|A\|$, independent of m. With a view towards using Lemma 5.4 we let $\Delta' = i({}_1\Delta_+)$. We let Ξ_m be Δ_1, the first inner terminal of Ω_w. We let Ξ' be the first inner terminal Δ'_1 of Ω'_w. The inner filled-in terminals of $[\Omega_+, S]$ are the same as two of the terminals of $]\Omega_+m, S[$. Therefore, by Lemma 4.2, we have $\Xi_m \to \Xi'$ barycentrically. By Lemma 5.4 we see that $A(\Delta_1) \to A'(\Delta'_1)$ barycentrically. But $A'(\Delta'_1)$ is one of the inner filled-in terminals of Ω_w. The same argument works for Δ_2. This completes the induction step. □

It follows from Lemma 5.5 that the limit sets Λ_{f_m} converge to $i(\Lambda)$ and in fact the maps $S^1 \to \Lambda_{f_m}$ converge pointwise to the map $S^1 \to \Lambda_f$. The action of ρ_{f_m} on Λ_{f_m} is determined by the embedding $S^1 \to \Lambda_{f_m}$ and by the action of $PSL_2(\boldsymbol{Z})$ on $\partial \boldsymbol{H}^2$. Hence the restriction of ρ_{f_m} to Λ_{f_m} converges to the restriction of $i \circ \rho_f \circ i^{-1}$ to Λ_f. This completes the proof of Theorem 1.4, except in the case when Γ_1 is the empty pattern.

The empty pattern To deal with the empty pattern we first define the *standard representation* of $PSL_2(\boldsymbol{Z})$ onto S^1. We think of the unit interval I_+ as the convex hull of the vectors $A = (1,0)$ and $C = (0,1)$ in $\boldsymbol{R}^2$. The midpoint of I_+ is the vector $B = (1/2, 1/2)$. Let $I_{1+} = \langle A \cup B \rangle$ and $I_{2+} = \langle B \cup C \rangle$. Here $\langle \cdot \rangle$ denotes the convex hull operation. Let I_- be another copy of $I+$. Let $S^1 = I_+ \cup I_-$. Let $\sigma_+ \in \mathrm{PL}(S^1)$ be the order 3 element whose action is given by the orbit $I_- \to I_{1+} \to I_{2+}$ and $A \to B \to C$. Let $\iota \in \mathrm{PL}(S^1)$ be the order 2 element whose action is given by the orbit $I_+ \to I_-$ and $A \to C$. The elements σ_+ and ι generate an action on S^1 which is topologically conjugate to the standard action of $PSL_2(\boldsymbol{Z})$ on $\partial \boldsymbol{H}^2$. To see this, we identify A, B, and C with the vertices of the ideal triangle t_+. Then σ_+ acts on S^1 just as the order 3 stabilizer of t_+ acts on $\partial \boldsymbol{H}^2$ and ι acts on S^1 just as the order 2 stabilizer of one of the edges of t_+ acts on $\partial \boldsymbol{H}^2$.

When Γ_1 is the empty pattern we are dealing with a sequence $\{f_m\}$ which converges to the 0–map. In this case, the same analysis as that given in Lemma 4.2 shows that the terminals of $]\Omega, S_m[$ converge to the line segments I_-, I_{1+} and I_{2+}. The same argument we give in Lemma 5.5 then shows that Λ_{f_m} converges to S^1 and the restriction of ρ_{f_m} to Λ_{f_m} converges to the standard representation. This completes our proof of Theorem 1.4.

References

[1] **L Bers**, *On Boundaries of Teichmuller Spaces and on Kleinian Groups I*, Annals of Math 91 570–600 (1970)

[2] **M Bonk**, **O Schramm**, *Embeddings of Gromov Hyperbolic Spaces*, Geom. Funct. Anal. 10 (2000) 266–306

[3] **F Gardiner**, *Teichmuller Theory and Quadratic Differentials*, Wiley Interscience (1987)

[4] **B Maskit**, *Kleinian Groups*, Springer–Verlag (1987)

[5] **J G Ratcliff**, *Foundations of Hyperbolic Manifolds*, Graduate Texts in Mathematics 149, Springer-Verlag (1994)

[6] **C Series**, *Geometric Methods of Symbolic Coding*, from: "Ergodic Theory, Symbolic Dynamics, and Hyperbolic Spaces" (Bedford, Keane and Series, editors) Oxford Science Publications (1991)

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

[8] **R E Schwartz**, *Circle Quotients and String Art*, Topology 41 (2001) 495–523

[9] **R E Schwartz**, *Pappus's Theorem and the Modular Group*, IHES Publications Mathematiques 78 (1993)