Automorphisms of hyperbolic surfaces

November 2013

0 Introduction

We want to classify the homeomorphisms $h : F \to F$ of a genus $g \geq 2$ surface F. Up to isotopy, that is, quotiented by the component of identity. Let's first look at the $g = 1$ case, where, just like for higher genus, homotopy classes of homeomorphisms are the same as isotopy classes, and the classes are determined by the induced map on $\pi_1(\mathbb{T}^2)$, which is an element of $SL_2(\mathbb{Z})$.

The classification is via the trace of the matrix A of h_* in some basis. If $\left|\mathrm{Tr}(A)\right| < 2$, then some algebra shows that h_* is periodic (in fact period dividing 12), and hence h^{12} is isotopic to identity.

If $\left|Tr(A)\right| = 2$, then there are integral eigenvectors (with eigenvalue ± 1), and hence closed curves that are preserved by h upto homotopy. We call such h reducible. The non-identity case is essentially a product of Dehn twists.

The last case, where $|\mathrm{Tr}\, A| > 2$, $A = \begin{pmatrix} \lambda & 0 \\ 0 & \frac{1}{\lambda} \end{pmatrix}$. Consider the real eigenvectors, which correspond to complementary slopes in $\mathbb{R}^2$. The translates of these transversely foliate $\mathbb{T}^2$, and h is dilation and contraction by the factor λ in these two directions. These maps are called Anosov. Starting with any closed curve C, $h^n(C)$ and $h^{-n}(C)$ "approach" these two foliations.

We are going to obtain an analogous classification for the higher genus case, except that there is no a priori explicit notion of trace. And the periodic case is a subset of the reducible case, which is defined to be somewhat more general.

1 The hyperbolic plane - $\mathbb{H}^2$

We will use the Poincare disk model for $\mathbb{H}^2$, with boundary denoted by S^1_∞. The group of isometries, $\mathcal{I}^+(\mathbb{H}^2) \simeq PSL_2(\mathbb{R})$ and is classified into elliptic, parabolic and hyperbolic according to fixed points.

2 Hyperbolic structures on surfaces

For $g \geq 2$, we can consider a hyperbolic structure on F, defined as follows:

Definition 1. A hyperbolic structure on a surface F is determined by an atlas of charts $\varphi_\alpha : U_\alpha \to \mathbb{H}^2$ such that the transition maps are restrictions of elements of $\mathcal{I}^+(\mathbb{H}^2)$. A surface with a hyperbolic structure is called a hyperbolic surface.

With some effort, one can show that, $F \simeq \mathbb{H}^2/\Gamma$ where $\pi_1(F) \simeq \Gamma < PSL_2(\mathbb{R})$ contains only hyperbolic isometries. In particular, this quotient map is a universal covering map. Geodesics in F are just images of geodesics in $\mathbb{H}^2$.

Definition 2. A geodesic in F is defined locally, that is, a pullback of the geodesics of $\mathbb{H}^2$ via the hyperbolic charts.

Definition 3. A closed curve in a surface (or any other space) is essential if it is not null-homotopic.

Our basic objects of consideration are non-trivial unions of loops:

Definition 4. A closed 1-submanifold of a surface F is a disjoint union of simple closed curves in F. An essential 1-submanifold is a closed 1-submanifold in which every component is essential and no two components are homotopic.

It's nice to know that they are uniquely represented by geodesics.

Lemma 2.1. *Every essential 1-submanifold C of a closed hyperbolic surface is isotopic to a unique geodesic 1-submanifold.*

And this gives us a good characterization of those h that should be called periodic.

Theorem 2.2. *Let $h : F \to F$ be an orientation preserving automorphism of a closed hyperbolic surface F. If for every closed essential curve C in F there is a $k > 0$ such that $h^k(C)$ is homotopic to C, then there exists an $n > 0$ such that h^n is isotopic to the identity.*

And we have some sort of non-compressibility result coming from Gauss-Bonnet, that will force finiteness in certain useful ways.

Theorem 2.3. *A complete hyperbolic surface F with finite area and geodesic boundary is homeomorphic to a compact surface less a finite set, and has area $-2\pi\chi_F + \pi\chi_{\partial F}$.*

3 Geodesic laminations

Now, we already know that we can take geodesics to represent closed curves without loss of generality. But we would like to take limits. Some foresight tells us that we should consider closed sets. And that essentially is what we call a lamination:

Definition 5. A (geodesic) lamination on F is a non-empty closed subset L of F which is a disjoint union of simple geodesics. The geodesics contained in L are called its leaves.

In the following assume $L = \cup_{x \in L} \gamma_x$ where γ_x is a geodesic passing through x, and γ_x, γ_y are either disjoint or coincident. We have not yet shown that this decomposition is unique. Now, the non-transverseness of this decompositions forces the direction of γ_x to vary continuously with x, as can be seen by charts.

Negative (non-zero) Euler characteristic tells us that $L \neq F$, since it's easy to turn a line field into a non vanishing vector field (possibly on some double-cover). Then the next picture tells us that L has no interior, and hence the "leaf" decomposition is unique.

Thus we have the following:

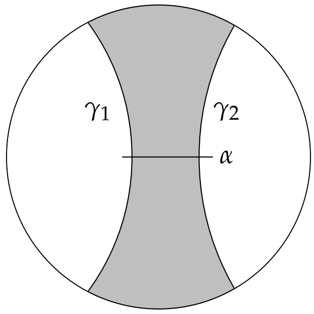

Figure 1: If α is a arc of geodesic contained in L and transverse to every γ_x it intersects, then L must contain the image of the whole shaded region which has to contain some fundamental domain for F. This contradicts $F \neq L$.

Lemma 3.1. *A geodesic lamination is nowhere dense, and has a unique leaf decomposition.*

The simplest way to construct laminations is take a bunch of disjoint geodesics, and take the closure:

Lemma 3.2. *The closure of a non-empty disjoint union L of geodesics is a lamination.*

We want to take limits, so:

Definition 6. $\Lambda(F)$ is the set of all geodesic laminations of F.

Topologize this by the Hausdorff metric (or with the hit-and-miss topology if you don't like metrics) on $K(F)$, the set of all compact subsets of F.

Theorem 3.3. $\Lambda(F)$ *is compact under the Hausdorff metric, if F is closed.*

The proof is somewhat complicated, since the limits need to be taken over $(x, \text{direction of } \gamma_x)$, that is, essentially in $K(PT(F))$, $PT(F)$ being the projective bundle. Now we want $h : F \to F$ to induce a map on $\Lambda(F)$. But before that, the following is standard:

Lemma 3.4. *Let $h : F_1 \to F_2$ be a homeomorphism of closed orientable hyperbolic surfaces, and let $\tilde{h} : \mathbb{H}^2 \to \mathbb{H}^2$ be a lift of h to the universal cover. Then $\tilde{h}$ has a unique extension to $\mathbb{H}^2 \cup S^1_\infty$.*

h is also determined upto homotopy by $\tilde{h}|_{S^1_\infty}$ in the sense that if $\tilde{h}, \tilde{k}$ agree on S^1_∞ then $h \approx k$, and conversely if h, k are homotopic, then they have lifts which agree on S^1_∞. This also generalizes to higher dimensions and as in Mostow rigidity, the premise can be weakened to h being a homotopy equivalence.

We use the same notation $\tilde{h}$ to denote both the extension, and the restriction to S^1_∞. Now this map can be used to send geodesics to geodesics, by looking at the endpoints (see fig. 2). In some sense, we identify the set of all geodesics on $\mathbb{H}^2$ with $\left\{ \{a, b\} \subset S^1 \;\middle|\; a \neq b \right\}$, and ones on F with this set quotiented by the action of Γ. Now since the map on S^1 is a homeomorphism, intersection properties of geodesics are preserved by this map. Thus we can extend to laminations via union. So, we have,

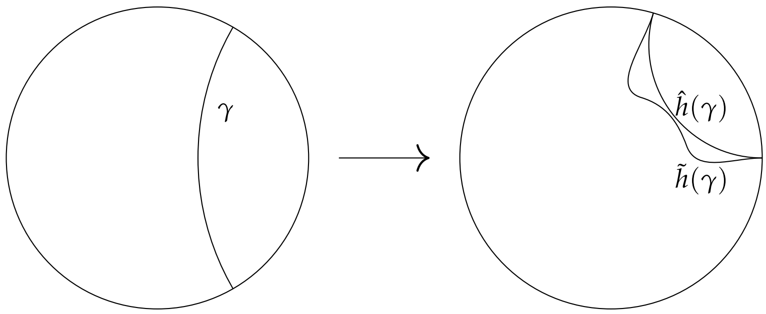

Figure 2: The map on geodesics

Theorem 3.5. *Any orientation-preserving homeomorphism $h : F_1 \to F_2$ of closed oriented hyperbolic surfaces induces a homeomorphism $\hat{h} : \Lambda(F_1) \to \Lambda(F_2)$. If h is homotopic to k then $\hat{h} = \hat{k}$. $\widehat{h_1 \circ h_2} = \widehat{h_1} \circ \widehat{h_2}$ whenever the compositions are defined.*

Now, almost obviously, $\hat{h}(\gamma)$ is homotopic to $h(\gamma)$. In fact, this also defines $\hat{h}$ without talking about the preimage, since there is a unique geodesic in the homotopy class of $h(\gamma)$. And further, the homotopy can be shifted from the level of images to the level of maps:

Corollary 3.6. *If $\hat{h}(C) = D$, then there is a k isotopic to h, such that $k(C) = D$ (and obviously, $\hat{k}(C) = D$ as well).*

4 Structures of geodesic laminations

Now, we want a lamination to be preserved by h, but h to not be "reducible", so we want to throw out closed leaves. As we'll see, throwing out closed leaves can also be done by throwing out "isolated" leaves. The "natural" definition for isolated leaves is as follows:

Definition 7. A leaf γ of L is isolated if for each $x \in \gamma$, there is a neighborhood U of x such that $(U, U \cap L)$ is homeomorphic to $(\text{disk}, \text{diameter})$.

By looking at the lifts, it's not very difficult to see that if the above holds for some $x \in \gamma$, then it holds for every $x \in \gamma$. And now,

Definition 8. $L' = L - \{\text{isolated leaves}\}$. Note that L' is closed. If non-empty, we call it the derived lamination of L.

If $L' = L$, we call L perfect. Now to develop the relations between closed and isolated leaves. If every leaf is isolated, then they are closed as well (and finitely many).

Lemma 4.1. *If L' is empty, then L is a finite union of simple closed geodesics and L is an isolated point in $\Lambda(F)$.*

To say more, we need to look at the complement sets. Since laminations are nowhere dense, these are open dense sets.

Definition 9. A component of $F \setminus L$ is a principal region for L.

We denote principal regions by U. Now the preimages of principal regions have contractible (even convex) components. This essentially follows from the boundary components being disjoint geodesics.

Lemma 4.2. *Let $\tilde{U}$ be a component of the preimage of U. Then $\overline{\tilde{U}}$ is a contractible hyperbolic surface with geodesic boundary.*

Now, since the universal cover of U is embedded in the universal cover of F, we have

Exercise 4.1. *$\pi_1(U) \to \pi_1(F)$ is injective.*

Now the boundary of the preimage corresponds to only some leaves of L, since it has to be isolated from one side. These we call boundary leaves:

Definition 10. A boundary leaf of a principal region U is a leaf γ of L such that for each $x \in \gamma$, there is an $\epsilon > 0$ such that $N_\epsilon(x) \cap U$ contains at least one component of $N_\epsilon \setminus \sigma$ where σ is the length 2ϵ segment of γ centered at x.

But we can recover the whole of L from the boundary leaves:

Lemma 4.3. *The union of all the boundary leaves of L is dense in L.*

Now, we can also take the closure of U in a different way:

Definition 11. Let $\Gamma_U = \left\{ g \in \Gamma \;\middle|\; g(\tilde{U}) = \tilde{U} \right\}$. Then, $U = \tilde{U}/\Gamma_U$. Define $V_U = \overline{\tilde{U}}/\Gamma_U$.

Now, for each principal region U, V_U is complete, hyperbolic, has geodesic boundary, and has interior homeomorphic to U. But we know such surfaces have area bounded below. Since F has finite area, we have:

Lemma 4.4. *L has only finitely many principal regions, each with only finitely many boundary leaves.*

Remark. Boundary leaves correspond to boundary components of G, but the correspondence is not necessarily 1-1.

Remark. The area of a principal region is $n\pi$ for some $n > 0$, so the maximum number of principal regions is $-2\chi_F = 4g - 4$.

Now say we can find a lamination without any closed leaves, then the principal regions have certain structure:

Lemma 4.5. *Let L have no closed leaves. Any principal region U is either isometric to a finite sided polygon with vertices on S^1_∞, or has a unique non-empty compact subset U_0 of U such that $U \setminus U_0$ is a finite disjoint union of interiors of crowns.*

Where a crown is a surface of a particular shape:

Definition 12. A crown is a complete hyperbolic surface W with finite area and geodesic boundary which is homeomorphic to $(S^1 \times [0,1]) \setminus A$ where A is a non-empty subset of $S^1 \times 1$. Note that ∂W has exactly one circle component.

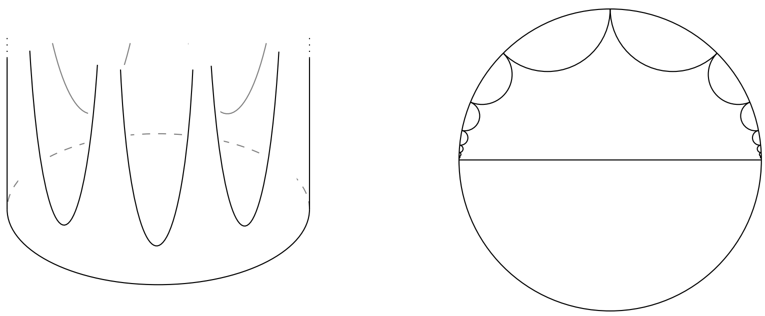

Figure 3: A crown and its universal cover

The proof of the classification of principal regions involves picking out certain elements of Γ that shift some boundary leaves of the preimage by some finitely many steps, so that one side of its axis, along with these leaves looks like the preimage of a crown. Then we just need to note that both this set, and its complement in $\tilde{U}$ is convex. All of this crucially uses the finiteness properties of the image. And then uniqueness again follows from the fact that all these sets in the preimage are forced to be convex, and contain certain leaves.

The compact subset is called the core.

Definition 13. U_0 in lemma 4.5 is the core of U. Each component of $U \setminus U_0$ is a crown of U. If U is a finite sided polygon, then we define the core of U to be empty.

We have some more finiteness results, even in the preimage, whenever there are no closed leaves.

Lemma 4.6. *Let L have no closed leaves. Then no point on S^1_∞ is an endpoint of infinitely many leaves of $\tilde{L}$.*

And this finiteness tells us that the leaves that are very close to closed leaves, are themselves isolated.

Lemma 4.7. *Any closed leaf C in L has a neighborhood N such that $L' \cap N \subset C$.*

So in particular, if $L' = L$, then L cannot have any closed leaves. The converse doesn't exactly hold, but we need just one more step:

Lemma 4.8. *For any lamination L, $L''' = L''$. If L has no closed leaves, then $L'' = L'$.*

This (and the next corollary) follows from the following technical result, which essentially says that minimal sublaminations of L' are its components.

Theorem 4.9. *If L_1 is a sublamination of L, then $L_1 \cap L'$ is a union of components of L'.*

The proof looks at the non-closed leaves of $L_1 \cap L$, which is by definition a lamination without closed leaves, and uses the finiteness and isolated leaf results above. The following is more or less the best we can say about laminations such that $L' = L$.

Corollary 4.10. *Every leaf of L is dense in L iff L is connected and $L' = L$ (or L is just one closed leaf).*

Now, let's see what we know about (connected) perfect laminations. They are nowhere dense, have no closed leaves, every leaf is dense. Further, only finitely many of the leaves are isolated on one side, and none on both sides. It should be clear by now that they are terribly difficult to draw. In fact, the intersection with any transverse arc is a cantor set. Fortunately, they are measure 0, so not as bad as fat cantor sets. The proof of this is some clever applications of Gauss-Bonnet and Poincare-Hopf, and is completely orthogonal to the rest of the arguments, so we will skip it.

We also have that all these notions related to laminations are preserved by the maps induced by homeomorphisms.

Lemma 4.11. *Let F_1, F_2 be closed orientable hyperbolic surfaces, let $h : F_1 \to F_2$ be a homeomorphism, and let L be a geodesic lamination in F_1. Then $\hat{h}(L)' = \hat{h}(L')$, and for any principal region U, there is a unique principal region $\hat{h}(U)$ for $\hat{h}(L)$ whose boundary leaves are $\hat{h}($boundary leaves of $L)$. If U has core U_0 then $\hat{h}(U)$ has core $\hat{h}(U_0)$ with frontier $\hat{h}(\partial U_0)$.*

The proof involves rewriting all the definition in terms of intersections of geodesics, which we already know are preserved by $\hat{h}$.

5 Surface automorphisms

Now we know enough about laminations to be able to classify the maps.

Definition 14. An automorphism h of a closed oriented surface F is called periodic if h^n is homotopic to the identity for some $n > 0$. h is reducible if h is homotopic to an automorphism which leaves invariant some essential closed 1-submanifold of F.

Remark. If h^n is homotopic to identity, then h is isotopic to g such that g^n is the identity. First proved by Fenchel and Nielsen.

Remark. If F has a hyperbolic structure, an automorphism h is reducible iff $\hat{h}(C) = C$ for some geodesic 1-submanifold C. For the forward direction, note that $\hat{h}$ depends only on the homotopy class, and for the reverse direction, use corollary 3.6.

Now it's not at all clear whether aperiodic irreducible homeomorphisms exist. However, the following gives us a sufficient condition on the map induced on homology, which is relatively easy to compute.

Lemma 5.1. *Let h induce $h_* : H_1(F) \to H_1(F)$, having matrix A with respect to some basis. If the characteristic polynomial $\chi_h(t)$ of A is irreducible in $\mathbb{Z}[t]$, has no root of unity as a zero, and is not a polynomial in t^n for any $n > 1$, then h is irreducible and aperiodic.*

The proof involves some basic topology, and linear algebra, and is orthogonal to the main discussion. But given this, it's not hard to check that certain classes of products of Dehn twists satisfy the above conditions, and hence are in fact aperiodic and irreducible.

Now that we have maps that are neither periodic nor reducible, let's see how much we can analyse them. First some notation.

Definition 15. If L_1 and L_2 are laminations then $L_1 \pitchfork L_2$ denotes the set of transverse intersections of L_1 with L_2.

From just aperiodic, we get a fixed lamination:

Lemma 5.2. *If h is a non-periodic automorphism, then $\hat{h}(L) = L$ for some $L \in \Lambda(F)$.*

The construction is important, so we include the whole proof.

Proof. By theorem 2.2 there is a simple closed geodesic $C \subset F$ such that $\hat{h}^n \neq C$ for all $n > 0$. It follows that $\hat{h}^m(C) \neq \hat{h}^n(C)$ for $n \neq m$.

Since $\Lambda(F)$ is compact (theorem 3.3), the sequence $\hat{h}^n(C)$ has a convergent subsequence, say $\hat{h}^{n_i}(C)$ converges to $K \in \Lambda(F)$. Since the geodesics $\hat{h}^{n_i}(C)$ are distinct, K is not an isolated point of $\Lambda(F)$. By lemma 4.1, K' is non-empty.

For a fixed r, theorem 3.5 implies that $\hat{h}^{n_i+r}$ converges to $\hat{h}^r(K)$. Note that

$$\left|\hat{h}^{n_i}(C) \cap \hat{h}^{n_i+r}(C)\right| = \left|C \cap \hat{h}^r(C)\right|$$

since composing with isotopies (via corollary 3.6) do not affect intersections. Let $N_r = \left|C \cap \hat{h}^r(C)\right|$ and suppose that $\left|K \pitchfork \hat{h}^r(K)\right| > N_r$. Choose disjoint neighborhoods of $N_r + 1$ points in $K \pitchfork \hat{h}^r(K)$. If n_i is sufficiently large, each contains at least one point of $\hat{h}^{n_i}(C) \cap \hat{h}^{n_i+r}(C)$. This is impossible, so $\left|K \pitchfork \hat{h}^r(K)\right| \leq N_r$ is finite. It follows that $K \pitchfork \hat{h}^r(K')$ is empty.

Then for any r, s, $\hat{h}^s(K') \pitchfork \hat{h}^r(K')$ is empty. So the union E of all $\hat{h}^r(K')$ is a non-empty disjoint union of simple geodesics. By lemma 3.2, $L = \bar{E}$ is a lamination. Clearly $\hat{h}(L) = L$. Observe, that $K \pitchfork L$ is empty, by the proof of lemma 3.2. □

As an easy corollary we have a seemingly independent result:

Corollary 5.3. *There is a lift $\tilde{k}$ of a strictly positive power of h such that the restriction of $\tilde{k}$ to S^1_∞ has at least two fixed points.*

But if h is reducible, then for all we know, we've obtained the isolated lamination which h fixes (upto homotopy). However, the following string of results improve on the constraints we have on L, given that h is irreducible. First of all, we have that L' and L are of the hard to draw type.

Lemma 5.4. *If h is an irreducible automorphism, and $\hat{h}(L) = L$ then each component of $F \setminus L'$ is contractible (and in fact a finite sided polygon), and each leaf of L is dense in L' (by which we mean that its closure contains L'). L has no closed leaves, and hence L' is perfect.*

The idea here is that the cores of principal regions have circle components in their boundaries. So if there is some non-empty core for some region, then the union of all these components would be a fixed 1-submanifold. Connectedness of L follows from some standard topological arguments (using Urysohn) given that the complement has contractible components.

Now we want to see what h does on the leaves of L. Looking ahead, the final property is that h expands along these leaves. Then by iteration of $\tilde{h}$, every point should converge to some endpoint of a leaf of $\tilde{L}$. On S^1, almost every point should. But then take two consecutive such contracting endpoints. This interval has to have leaves approaching the endpoints. Given this, the following definition and lemma are somewhat natural.

Definition 16. Let L be a geodesic lamination and $\tilde{L}$ be its preimage. A stable interval for L is a closed interval $I \subset S^1_\infty$ such that for any two points $P, Q \in \mathring{I}$, there is a leaf $\tilde{\gamma} \subset \tilde{L}$ whose endpoints separate P and Q from ∂I.

Lemma 5.5. *If $h : F \to F$ is non-periodic, irreducible, then $\hat{h}(L) = L$ for some L with the following property. If a lift $\tilde{k}$ of a strictly positive power of h maps a stable interval I onto itself, then the restriction of $\tilde{k}$ to I has a fixed point $Z \in \mathring{I}$ such that for all $P \in I \setminus Z$, $\tilde{k}^n(P)$ converges to a point in ∂I.*

The proof is basically a lot of case checking, along with with repeated uses of the fact that two geodesics intersect iff their endpoints alternate.

Remark. Since $\tilde{k}$ is orientation preserving and $\tilde{k}(I) = I$, the endpoints of I are fixed by $\tilde{k}$. The endpoints are contracting fixed points of the restriction of $\tilde{k}$ to I and Z is an expanding fixed point. Note that the restriction of $\tilde{k}$ to I has no other fixed points.

But after having shown this, we can eliminate all other types of fixed points, and then compactness tells us that there are only finitely many, as in the following theorem:

Theorem 5.6. *Let h be non-periodic, irreducible. Any lift of a strictly positive power of h has finitely many fixed points on S^1_∞, alternately contracting and expanding. There is a unique perfect lamination L^S, invariant under $\hat{h}$, such that $\tilde{L}^S$ contains the geodesics joining consecutive contracting fixed points of any lift of a strictly positive power of h. Every leaf of L^S is dense in L^S.*

Again, more case checking, and using the fact that any point on S^1 is the endpoint of finitely many leaves. Uniqueness essentially follows from the fact that once some power of h has at least two fixed points, the geodesic joining (some two of) them is both contained and dense in L^S, so determines it.

Definition 17. L^S is the stable lamination of h. The stable lamination L^U of h^{-1} is the unstable lamination of h.

A couple of results follow from this, the first basically says that if h is aperiodic irreducible, then all powers of h are irreducible.

Lemma 5.7. *If h is non-periodic, irreducible and C is an essential closed curve in F, then $h^m(C)$ is homotopic to $h^n(C)$ iff $m = n$.*

And the next strengthens the uniqueness L^S a bit.

Theorem 5.8. *There are only finitely many laminations K_i containing L^S. There is an integer $m > 0$ such that for any simple closed geodesic C, the sequence $\hat{h}^{mn}(C)$ converges to some K_i.*

And one more topological result tells us that the homotopy can be between L^S and $h(L^S)$ can be used to isotope h so that L^S is fixed under the homeomorphism:

Theorem 5.9. *If h is aperiodic, irreducible, then there is an h' isotopic to h such that $h'(L^S) = L^S$, $h'(L^U) = L^U$.*

The proof proceeds in three steps. First we observe that $\tilde{h}$ restricts to a bijection of $\tilde{L}^S \cap \tilde{L}^U$. Then we extend this linearly over $\tilde{L}^S \cup \tilde{L}^U$, and then to the complement $F \setminus (L^S \cup L^I)$ (cofinitely many components of which are rectangles).

Since none of this changes the map on S^1, the resulting map is homotopic to h.

6 Pseudo-Anosov automorphisms

Now, given an aperiodic irreducible h, we've analysed what it does on a nowhere dense subset of F. To obtain the full result, we need to fill the lamination to a foliation. In fact, we want a pair of transverse foliations, which we define below:

Definition 18. Two singular foliations $\mathcal{F}^S, \mathcal{F}^U$ are transverse if they have the same singular sets, the leaves are transverse everywhere else, and at the singular points, the separatrices alternate.

Now we use the fact that we remarked before, that the intersection of L^S (or L^U) with any transverse arc is a Cantor set. Now let $C \subset I$ be a cantor set, and for every $]a, b[\subset I \setminus C$, let $[a, b]$ be equivalent. Then it's not too hard to show that the quotient map is a homotopy equivalence from $I \to I/ \sim \simeq I$. Extending this to two dimensions, using L^S and L^U, consider the equivalence relation generated by:

1. If A is a component of $F \setminus (L^S \cup L^U)$, then $\overline{A}$ is equivalent.
2. If A is a component of $L^S \setminus L^U$, then $\overline{A}$ is equivalent.
3. If A is a component of $L^U \setminus L^S$, then $\overline{A}$ is equivalent.

Then by similar arguments, the quotient map π a homotopy equivalence. For genus g surfaces, this implies that π is homotopic to a homeomorphism $\Theta : F \to F/ \ \simeq F$. It is not too hard to check that pushing by π and pulling back by Θ results in the laminations L^S and L^U producing transverse foliations denoted $\mathcal{F}^S$ and $\mathcal{F}^U$. So we have the following lemma:

Lemma 6.1. *If h is aperiodic, irreducible, then h is isotopic to a map h_* such that there exist transverse singular foliations $\mathcal{F}^S, \mathcal{F}^U$ which are preserved by h_*.*

We push forward and pull back h' obtained in theorem 5.9 to get h_*. Since we are anyway interested in isotopy classes, fix the $\mathcal{F}^S, \mathcal{F}^U$ constructed above, and assume that h itself preserves them. We can also make the following observations about the foliations:

Lemma 6.2. *The foliations $\mathcal{F}^S, \mathcal{F}^U$ have no closed leaves. Each separatrix has only one singular point and is dense in the surface. If σ is a separatrix of $\mathcal{F}^U$ beginning at a singularity s, then there is an integer $m > 0$ such that $h^m(\sigma) \subset \sigma$ and for all $x \in \sigma$, $h^m(x)$ is contained in the open subinterval (s, x). A similar result holds for $\mathcal{F}^S$ and h^{-1}.*

Now we want to construct what are called Markov Partitions on F.

Lemma 6.3. *There are closed sets $A, B \subset F$ with the following properties:*

1. *A is a union of closed arcs, each contained in a separatrix of $\mathcal{F}^S$.*
2. *B is a union of closed arcs, each contained in a separatrix of $\mathcal{F}^U$.*
3. *Each component of $F \setminus (A \cup B)$ is homeomorphic to $]0, 1[\times]0, 1[$ with the leaves of $\mathcal{F}^S, \mathcal{F}^U$ being carried to horizontals and verticals respectively.*
4. *$h^{-1}(A) \subset A$.*
5. *$h(B) \subset B$.*

The construction is just a clever choice of picking segments on each separatrix, so that every closed arc contains exactly one singularity. Once such A, B have been constructed, we get the following corollary, which exactly describes the existence of a Markov partition. But first we define rectangles on F which are basically ones with "sides parallel to" the foliations.

Definition 19. A rectangle on F is a map $\rho : I \times I \to F$ such that ρ is an embedding on the interior, $\rho(\text{point} \times I)$ is contained in a leaf of $\mathcal{F}^U$ and $\rho(I \times \text{point})$ is contained in a leaf of $\mathcal{F}^S$. We confuse the map ρ with its image R. Also, denote $\rho(\partial I \times I)$ by $\partial^U R$ and $\rho(I \times \partial I)$ by $\partial^S R$.

Corollary 6.4. *F can be partitioned into finitely many rectangles R_i such that*

1. *For $i \neq j$, R_i and R_j have disjoint interiors.*
2. *$h(\cup \partial^U R_i) \subset \cup \partial^U R_i$.*
3. *$h^{-1}(\cup \partial^S R_i) \subset \cup \partial^S R_i$.*

So $h(R_i) \cap R_j$ is finitely many subrectangles stretched across R_j horizontally. Denote the number of these subrectangles by a_{ij}.

Now, the final result we want to prove, the analogue of the Anosov automorphisms for the $g = 1$ case, is the following:

Theorem 6.5. *If h is an aperiodic irreducible automorphism that preserves a pair of transverse singular foliations $\mathcal{F}^S$ and $\mathcal{F}^U$, then there exist transverse measures μ^S and μ^U respectively and some $\lambda > 1$ such that $h(\mu^S) = \lambda \mu_S$, $h(\mu^U) = \frac{1}{\lambda}\mu^U$.*

where we make use of the following definition:

Definition 20. A transverse measure μ on a foliation $\mathcal{F}$ is a measure on arcs transverse to $\mathcal{F}$ and avoiding its singular set, which is preserved by homotopies along $\mathcal{F}$. For a homeomorphism h preserving $\mathcal{F}$, we define $h(\mu)(h(\alpha)) = \mu(\alpha)$ for any transverse arc α.

Such automorphisms are called "pseudo-Anosov", again in analogy to the case of the torus. To prove the theorem, we just need to construct the transverse measures μ^S, μ^U. Suppose for now that such measures do exist, assigning the rectangle R_i a height y_i. Then using the matrix A with entries a_{ij} defined before, we must have,

$$y_j = \sum a_{ij} \lambda^{-1} y_i$$

since each subrectangle of $h(R_i) \cap R_j$ must have height (measured by μ^S, since transverse to $\mathcal{F}^S$) $\lambda^{-1} y_i$. In other words, (y_i) is an eigenvector of A with eigenvalue λ. So a necessary condition is that A must have an eigenvalue > 1 and a corresponding eigenvector with non-negative entries. We in fact have something stronger:

Lemma 6.6. *A has an eigenvalue $\lambda > 1$ with eigenvector $y = (y_i)$ such that each $y_i > 0$.*

An easy variation of Perron-Frobenius produces a eigenvector with non-negative entries, and each separatrix being dense (this is some sense ensures sufficient mixing) contradicts $y_i = 0$ and also shows that $\lambda > 1$.

Now, for a transverse arc α, letting $u_{i,m}$ be the number of components of $\alpha \cap h^m(R_i)$, the limit

$$\lim_{m\to\infty} \sum_i \lambda^{-m} y_i u_{i,m}$$

exists (consecutive terms differ by $O(\lambda^{-m})$), and works as $\mu^S(\alpha)$.

By an analogous construction, we get μ^U with $h^{-1}(\mu^U) = \bar{\lambda}\mu^U$. Considering the product measure $\mu^S \times \mu^U$ on F, we get $\lambda = \bar{\lambda}$, which finishes the proof.