Algebraic & Geometric Topology 9 (2009) 1177–1199

The first cohomology of the mapping class group with coefficients in algebraic functions on the $\mathrm{SL}_2(\mathbf{C})$ moduli space

JØRGEN ELLEGAARD ANDERSEN
RASMUS VILLEMOES

Consider a compact surface of genus at least two. We prove that the first cohomology group of the mapping class group with coefficients in the space of algebraic functions on the $\mathrm{SL}_2(\mathbf{C})$ moduli space vanishes. In the genus one case, this cohomology group is infinite dimensional.

20J06; 57M07, 57M60

1 Introduction

Let Σ be a compact surface, possibly with boundary, of genus at least 2, and let $\mathcal{M} = \mathcal{M}_{\mathrm{SL}_2(\mathbf{C})}$ denote the moduli space of flat $\mathrm{SL}_2(\mathbf{C})$ connections over Σ. Since $\mathcal{M}$ may be identified with the space of $\mathrm{SL}_2(\mathbf{C})$ representations of the fundamental group of Σ modulo conjugation, $\mathcal{M}$ has the structure of an affine algebraic variety. The mapping class group Γ acts on $\mathcal{M}$ and hence on the space $\mathcal{O} = \mathcal{O}(\mathcal{M})$ of algebraic functions on $\mathcal{M}$, making $\mathcal{O}$ a module over Γ. The purpose of the present paper is to prove:

Theorem 1.1 *The first cohomology group $H^1(\Gamma, \mathcal{O})$ vanishes.*

The proof relies crucially on the Γ–equivariant identification of $\mathcal{O}$ with another vector space on which the action of Γ is more transparent. Based on Goldman's idea of using curves in the surface to represent functions on the moduli space, Bullock, Frohman and Kania-Bartoszyńska [5] (see also Skovberg [11]) proved that $\mathcal{O}$ is Γ–equivariantly isomorphic to the complex vector space spanned by the set of multicurves on Σ. This allows one to decompose $\mathcal{O}$ into smaller Γ–modules indexed by the mapping class group orbits of multicurves.

Letting $\mathcal{O}^*$ denote the algebraic dual of $\mathcal{O}$, the above identification of $\mathcal{O}$ with the complex vector space spanned by the set of multicurves, an application of Shapiro's

Published: 10 June 2009 DOI: 10.2140/agt.2009.9.1177

Lemma gives a description of $H^1(\Gamma, \mathcal{O}^*)$ in terms of stabilizers of multicurves. Using the set of multicurves as a basis also yields an inclusion map $\iota\colon \mathcal{O} \to \mathcal{O}^*$, which induces a map on cohomology, $\iota_*\colon H^1(\Gamma, \mathcal{O}) \to H^1(\Gamma, \mathcal{O}^*)$. The proof of Theorem 1.1 essentially consists of two steps. The first uses the description of the target to prove that ι_* is the zero map; this is the contents of Proposition 6.1. The second step consists of proving that ι_* is injective. This is done by introducing the notion of almost invariant colorings of a G–set, and then proving that no such almost invariant colorings exist in the case of a mapping class group orbit of a multicurve.

By a result of Goldman, the restriction map from $\mathcal{O}$ to continuous functions on the SU(2) moduli space is injective. This follows from the fact that the SU(2) moduli space is a real slice inside the $\mathrm{SL}_2(\mathbf{C})$ moduli space and the fact that the latter is irreducible. For further details, see the proof of Theorem 1.1 in Goldman [8]. See also Charles and Marche [6] for an alternative argument for this fact. Hence we get that $\mathcal{O}$ is a mapping class group invariant submodule of the continuous functions on the SU(2) moduli space, whose first cohomology group of course vanishes by our main theorem.

This paper is organized as follows. In the next section, we describe the background for this result. In Section 3, we describe the isomorphism of the ring of algebraic functions with the complex vector space spanned by multicurves and how this gives a splitting of the cohomology group into components indexed by mapping class group orbits. Section 4 introduces the dual coefficient module, and it is described how a standard result from group cohomology allows us to compute the cohomology with these coefficients in terms of stabilizers of multicurves. In Section 5 we recall certain standard facts about the mapping class group, relations between Dehn twists and their action on multicurves. The first part of the proof of Theorem 1.1 is done in Section 6. The notion of almost invariant colorings is introduced in Section 7, where we also prove that a mapping class group orbit of a multicurve admits no such (nontrivial) coloring. Combining these results, we finally prove Theorem 1.1 in Section 8.

In Section 9, we treat the case of a closed genus 1 surface. In that case, the cohomology group $H^1(\Gamma, \mathcal{O})$ turns out to be of infinite dimension. This is the only place where we consider $g = 1$; in all other sections the genus is assumed to be a least 2.

We thank the referee for valuable comments.

2 Motivation

The motivation for studying the first cohomology group of the mapping class group with coefficients in the space of functions on the moduli space came from [1]. In that paper, the first author studied deformation quantizations, or star products, of the Poisson algebra of smooth functions on the moduli space $\mathcal{M}_G$ of flat G–connections, where $G = \mathrm{SU}(n)$. The construction uses Toeplitz operator techniques and produces a family of star products parametrized by Teichmüller space. One step towards turning this family into one mapping class group invariant star product is taken by the first author and Gammelgaard in [2], where the existence to first order of a mapping class group equivariant formal trivialization of the so-called formal Hitchin connection is proved. The extent to which a mapping class group invariant star product is unique is also interesting. One result [1, Proposition 6] is that, provided the cohomology group $H^1(\Gamma, C^\infty(\mathcal{M}_G))$ vanishes, one may find a Γ–invariant equivalence between any two equivalent star products. Since it is easy to see that the only Γ–invariant equivalences are the multiples of the identity, this immediately implies that within each equivalence class of star products, there is at most one Γ–invariant star product.

Theorem 1.1 is clearly a step towards verifying the assumption above in the case of $G = \mathrm{SU}(2)$, since the $\mathrm{SU}(2)$–moduli space is included in the $\mathrm{SL}_2(\mathbf{C})$ moduli space.

3 Splitting the coefficient module

A *multicurve* is the isotopy class of a finite collection of pairwise disjoint, simple closed curves on Σ. Let B denote the set of multicurves on Σ, and let $\mathcal{B} = \mathcal{B}(\Sigma) = \mathbf{C}B$ denote the complex vector space spanned by B. In Skovberg [11] one finds a complete proof of the following:

Theorem 3.1 *There exists a Γ–equivariant isomorphism $\nu\colon \mathcal{B} \to \mathcal{O}$.*

If $D = \bigsqcup_j \gamma_j$ is the disjoint union of simple closed curves γ_j, $\nu(D)$ is simply $\prod_j -f_{\vec{\gamma}_j}$, where $\vec{\gamma}_j$ denotes any of the oriented versions of γ_j, and $f_{\vec{\gamma}_j}$ is Goldman's holonomy function on the moduli space.

Theorem 3.1 allows us to split $\mathcal{O}$ according to the mapping class group orbits of multicurves. More precisely, for a multicurve D, let $M_D = \mathbf{C}(\Gamma D)$ denote the complex vector space spanned by the Γ–orbit through D. Then we have a decomposition as Γ–modules

$$\mathcal{O} \cong \mathcal{B} \cong \bigoplus_D M_D \tag{1}$$

where the sum is over a set of representatives of the mapping class group orbits of multicurves. Since the mapping class group is known to be finitely generated, this induces a corresponding decomposition of the cohomology

$$H^1(\Gamma,\mathcal{B}) \cong \bigoplus_D H^1(\Gamma, M_D). \tag{2}$$

Hence it suffices to show that each summand on the right-hand side of (2) vanishes in order to prove Theorem 1.1.

4 The dual module

It turns out to ease the computation of $H^1(\Gamma, M_D)$ if one introduces a larger module. Let $\mathcal{B}^*$ denote the algebraic dual of $\mathcal{B}$. Using the set of multicurves as a basis, there is a Γ–equivariant inclusion $\mathcal{B} \to \mathcal{B}^*$. In fact, we may identify $\mathcal{B}^*$ with the space $\mathrm{Map}(B, \mathbf{C})$ of all formal linear combinations of multicurves. There is a decomposition of $\mathcal{B}^*$ similar to (1) into a direct product of Γ–modules,

$$\mathcal{B}^* \cong \prod_D \widehat{M}_D, \tag{3}$$

where $\widehat{M}_D = \mathrm{Map}(\Gamma D, \mathbf{C})$ denotes the set of all formal linear combinations of elements of the orbit through D, and the product is over the same set of representatives as in (1).

The Γ–equivariant inclusion $\iota\colon M_D \to \widehat{M}_D$ induces a long exact sequence in cohomology, the first part of which is

$$\begin{aligned}0 \to H^0(\Gamma, M_D) \to H^0(\Gamma, \widehat{M}_D) \to H^0(\Gamma, \widehat{M}_D/M_D)\\ \to H^1(\Gamma, M_D) \to H^1(\Gamma, \widehat{M}_D).\end{aligned} \tag{4}$$

In [3], we computed $H^1(\Gamma, \widehat{M}_D)$ for any multicurve D, and showed that for any surface there exists a multicurve such that $H^1(\Gamma, \widehat{M}_D)$ is nonzero.

We need the description of $H^1(\Gamma, \widehat{M}_D)$ given in [3], so let us recall the most important facts. Let $\Gamma_D \subseteq \Gamma$ denote the stabilizer of D in Γ (permutation of the components of D are allowed). Then the Γ–equivariant identification of the set Γ/Γ_D of left cosets with the orbit ΓD induces an isomorphism of $\widehat{M}_D = \mathrm{Map}(\Gamma D, \mathbf{C})$ with the space $\mathrm{Hom}_{\mathbf{Z}\Gamma_D}(\mathbf{Z}\Gamma, \mathbf{C})$ of $\mathbf{Z}\Gamma_D$–homomorphisms $\mathbf{Z}\Gamma \to \mathbf{C}$.

This Γ–module is also known as the coinduced module $\mathrm{Coind}_{\Gamma_D}^{\Gamma}\mathbf{C}$, and Shapiro's Lemma (see Brown [4]) yields an isomorphism

$$H^1(\Gamma, \widehat{M}_D) = H^1(\Gamma, \mathrm{Coind}_{\Gamma_D}^{\Gamma}\mathbf{C}) \cong H^1(\Gamma_D, \mathbf{C}) \tag{5}$$

where $\mathbf{C}$ is a trivial Γ_D–module. Hence $H^1(\Gamma, \widehat{M}_D)$ is simply the space of homomorphisms from (the abelianization of) Γ_D to $\mathbf{C}$.

Explicitly, the isomorphism (5) is given as follows: An element of $H^1(\Gamma, \widehat{M}_D)$ is represented by a cocycle $u\colon \Gamma \to \widehat{M}_D$, which can also be considered as a map $u\colon \Gamma \times \Gamma D \to \mathbf{C}$. Restricting to the subset $\Gamma_D \times \{D\} \equiv \Gamma_D$ we obtain a map $u_|\colon \Gamma_D \to \mathbf{C}$, which is easily seen to be a homomorphism. In other words, $u_|(g)$ is given by picking out the coefficient of D in $u(g)$.

4.1 Isomorphisms of modules

If D is a multicurve, let D^n denote the multicurve obtained from D by replacing each component by n parallel copies. Clearly, there are Γ–isomorphisms $M_D \to M_{D^n}$ and $\widehat{M}_D \to \widehat{M}_{D^n}$. Also, if γ is a simple closed curve parallel to a boundary component of Σ, we have Γ–isomorphisms $M_D \to M_{D\cup\gamma}$ and $\widehat{M}_D \to \widehat{M}_{D\cup\gamma}$. These observations imply that we may without loss of generality only consider multicurves without boundary parallel components, and satisfying that the multiplicities of the different components are relatively prime. Using nonstandard terminology, such a multicurve will be called *reduced*.

5 Dehn twists and multicurves

Before starting actual computations leading to a proof of Theorem 1.1, we need to record a few facts regarding Dehn twists, multicurves and the modules M_D, $\widehat{M}_D$.

5.1 Presentations and relations

It is well-known that the mapping class group is generated by Dehn twists. In fact, there exists a finite set of curves such that the Dehn twists on these curves generate Γ. Furthermore, one may choose these curves so that any pair of them intersect in at most two points (see Gervais [7]); if Σ has at most one boundary component, the curves may be chosen so that each pair intersect in at most one point (see Wajnryb [12]).

For later use, we mention a few relations between Dehn twists.

Lemma 5.1 *Dehn twists on disjoint curves commute.*

Lemma 5.2 *If α and β are simple closed curves intersecting transversely in a single point, the associated Dehn twists are* braided. *That is, $\tau_\alpha\tau_\beta\tau_\alpha = \tau_\beta\tau_\alpha\tau_\beta$.*

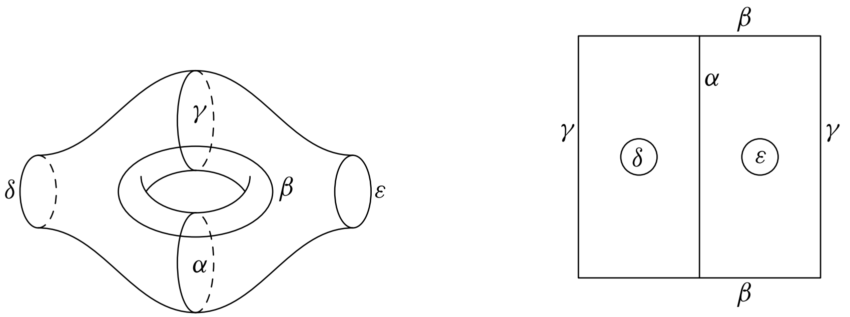

(a) A two-holed torus

(b) A more schematic picture. The opposite sides of the square are identified in the usual way to obtain a torus.

Figure 1: The chain relation

Lemma 5.3 (Chain relation) *Let α, β and γ be simple closed curves in a two-holed torus as in Figure 1, and let δ, ε denote curves parallel to the boundary components of the torus. Then $(\tau_\alpha\tau_\beta\tau_\gamma)^4 = \tau_\delta\tau_\varepsilon$.*

Lemma 5.4 *When $g \geq 2$, a twist on a separating curve can be written as a product of twists on nonseparating curves.*

Corollary 5.5 *The mapping class group is generated by a finite set of twists in nonseparating curves (though we may not necessarily choose this set so that each pair of curves intersect in at most two points).*

5.2 The action of twists on multicurves

There is a simple way to parametrize the set of all multicurves which was found by Dehn. For details, we refer to Penner and Harer [9]. Essentially one cuts the surface into pairs of pants using $3g+r-3$ simple closed curves γ_k, and then for each pants curve γ_k one records the geometric intersection number $m_k(D) = i(\gamma_k, D)$ (which is a nonnegative integer) and a "twisting number" $t_k(D)$, which can be any integer. This defines a $(6g+2r-6)$–tuple of integers $(m_1(D), t_1(D), \dots, m_{3g+r-3}(D), t_{3g+r-3}(D))$ (satisfying certain conditions), and, conversely, from any such tuple satisfying these conditions one may construct a multicurve.

The important fact is that in this parametrization, the action of the twist in the curve γ_k on a multicurve D is given by

$$t_k(\tau_{\gamma_k}^{\pm 1} D) = t_k(D) \pm m_k(D), \tag{6}$$

all other coordinates being unchanged. The formula (6) is intuitive in the sense that it says that for each time D intersects γ_k essentially, the action of τ_{γ_k} on D adds 1 to the twisting number of D with respect to γ_k. This can be used to prove a number of important facts.

Lemma 5.6 *Let γ be a simple closed curve and D a multicurve. Then the following are equivalent:*

(1) *The twist τ_γ acts trivially on D.*

(2) *The twist τ_γ acts trivially on each component of D.*

(3) *The geometric intersection number between γ and D is zero.*

(4) *One may realize γ and D disjointly.*

Conversely, if τ_γ acts nontrivially on D, all the multicurves $\tau_\gamma^n D$, $n \in \mathbf{Z}$, are distinct.

Proof All of the above assertions can be proved from (6) by letting γ be part of a pants decomposition of the surface. This is clearly possible if γ is nonseparating, while if γ is separating, observe that both connected components resulting from cutting along γ must have negative Euler characteristic (otherwise γ would be trivial or parallel to a boundary component, in which case the twist on γ clearly acts trivially on D). □

To find a twist acting nontrivially on a multicurve, we need only find a curve which has positive geometric intersection number with the multicurve. This is possible if and only if the multicurve has a component which is not parallel to a boundary component of Σ.

On a surface with negative Euler characteristic, there exist complete hyperbolic metrics of constant negative curvature. Within each free homotopy class of simple closed curves, there is a unique geodesic representative with respect to such a metric. If a and b are the geodesic representatives of distinct homotopy classes α, β, then a and b realizes the geometric intersection number between α and β, ie $\# a \cap b = i(\alpha, \beta)$.

6 The map to the cohomology with dual coefficients

Let D be a reduced multicurve. The purpose of the section is to prove:

Proposition 6.1 *The map $\iota_*\colon H^1(\Gamma, M_D) \to H^1(\Gamma, \widehat{M}_D)$ is zero.*

The proof uses the description of $H^1(\Gamma, \widehat{M}_D)$ as $\mathrm{Hom}(\Gamma_D, \mathbf{C})$ given at the end of Section 4. Let $u\colon \Gamma \to M_D$ be a cocycle. Since Γ is generated by Dehn twists, it is natural to study to which extent $u(\tau_\alpha)$ can contain nonzero terms on which τ_α acts trivially for simple closed curves α.

Lemma 6.2 *Let α be a simple closed curve on Σ, and let $E \in \Gamma D$ be a multicurve such that $\tau_\alpha E = E$. Assume that E contains at least one component which is not a parallel copy of α. Then the coefficient of E in $u(\tau_\alpha)$ is zero.*

Proof Let ε be a component of E which is not parallel to α. Then since every component of E is disjoint from α, and since we assumed that D (and hence E) is a reduced multicurve, ε is not parallel to a boundary component of the (possibly disconnected) surface Σ_α obtained by cutting Σ along α. Hence we may find a curve β disjoint from α such that $\tau_\beta\varepsilon \neq \varepsilon$ and thus $\tau_\beta E \neq E$. Then τ_α and τ_β commute, and $u(\tau_\alpha\tau_\beta) = u(\tau_\beta\tau_\alpha)$. Using the cocycle condition this becomes

$$u(\tau_\alpha) + \tau_\alpha u(\tau_\beta) = u(\tau_\beta) + \tau_\beta u(\tau_\alpha),$$

which we may rewrite as

$$(1 - \tau_\beta)\cdot u(\tau_\alpha) = (1 - \tau_\alpha)\cdot u(\tau_\beta). \tag{7}$$

Now since $\tau_\alpha E = E$, the coefficient of E on the right-hand side of (7) is clearly 0. Assuming that $u(\tau_\alpha)$ contains some nonzero term xE, (7) then implies that it must also contain the term $x\tau_\beta^{-1}E$. But since τ_α and τ_β commute, τ_α also acts trivially on $\tau_\beta^{-1}E$, so we may repeat the above argument with $\tau_\beta^{-1}E$ instead of E and conclude that $u(\tau_\alpha)$ then also contains the term $x\tau_\beta^{-2}E$. Continuing in this way, $u(\tau_\alpha)$ contains infinitely many nonzero terms (since the multicurves $\tau_\beta^n E$ are all distinct), which is impossible since we assumed that u took values in M_D. □

In other words, τ_α acts nontrivially on "most" of the nonzero terms occurring in $u(\tau_\alpha)$; the possible exception is when D consists of a single component and the curve α is in the orbit of D (eg if D and α are nonseparating curves). But this possibility is easily ruled out.

Proposition 6.3 *Let ε be any simple closed curve. Then τ_ε acts nontrivially on any nonzero term occurring in $u(\tau_\varepsilon)$.*

Proof By the previous lemma, we only need to prove that $u(\tau_\varepsilon)$ does not contain some nonzero term $x\varepsilon$, where ε is considered as a 1–component multicurve. To see this, observe that any curve ε can be realized as the ε occurring in the chain relation (Lemma 5.3); that is, there exists a genus 1 subsurface of Σ with two boundary components, one of which is ε: If ε is separating, one of the connected components obtained by cutting along ε has genus ≥ 1, and we may if necessary choose δ to be null-homotopic. If ε is nonseparating, it is always possible to find a δ such that the two curves together bound a genus 1 subsurface.

Applying the cocycle u to the chain relation, we obtain

$$u((\tau_\alpha \tau_\beta \tau_\gamma)^4) = u(\tau_\delta) + \tau_\delta u(\tau_\varepsilon).$$

But the left-hand side can be expanded (via the cocycle condition) to a sum of various actions of $\tau_\alpha, \tau_\beta, \tau_\gamma$ on the values of u on these twists; since they all act trivially on ε the coefficient of ε on the left-hand side is 0 by Lemma 6.2. Similarly, δ acts trivially on ε, so also the coefficient of ε in $u(\tau_\delta)$ is 0, and hence the coefficient of ε in $u(\tau_\varepsilon)$ is 0. □

Proof of Proposition 6.1 Let $u\colon \Gamma \to M_D$ be a cocycle. By the isomorphism (5) it suffices to prove the following: For any diffeomorphism $f \in \Gamma_D$ fixing the multicurve D, the coefficient of D in $u(f)$ is zero.

Since $u_|\colon \Gamma_D \to \mathbf{C}$ is a homomorphism, we may consider any power of f. Choose n sufficiently large so that f^n fixes each component of D and each side of each component. Then f^n may be realized as a diffeomorphism of the surface Σ_D obtained by cutting Σ along D. This implies that f^n can be written as a product of Dehn twists in curves not intersecting D. Hence, by Proposition 6.3, the coefficient of D in $u(f^n)$ is zero, and so is the coefficient of D in $u(f)$. □

7 Almost invariant colorings

Let G be a group and X an infinite set on which G acts. We define a *coloring* (or C–coloring) of X to be any map $c\colon X \to C$ into some set C of "colors". We will use the following terminology:

- A coloring c is *invariant* if $c(gx) = c(x)$ for each $g \in G$ and $x \in X$.

- A coloring is *almost invariant* if, for each $g \in G$, the identity $c(x) = c(gx)$ fails for only finitely many $x \in X$.
- Two colorings are *equivalent* if they assign different colors to only finitely many elements of X; this is clearly an equivalence relation on the set of C–colorings.
- A coloring is *trivial* if it is equivalent to a monochromatic (constant) coloring.

We will only deal with the case where the action of G is transitive. Then clearly the only invariant colorings are the constant ones, and hence we are only interested in studying the question of existence of almost invariant colorings. If two colorings are equivalent and one is almost invariant, so is the other, which explains the above definition of a trivial coloring. If one wants to classify all almost invariant colorings, this can clearly not be done better than up to the equivalence defined above.

A *simplification* of c is a coloring obtained by postcomposing c with some map $i\colon C \to C'$ (one "identifies" some of the colors). Clearly a simplification of an almost invariant coloring is almost invariant. Now, if there exists an almost invariant, nontrivial C–coloring c, there also exists an almost invariant coloring where exactly two colors are used. To see this, partition C into $C_0 \sqcup C_1$ such that $c^{-1}(C_k)$, $k = 0, 1$, are both infinite, and define a $\{0, 1\}$–coloring by composing c with the map $i\colon C \to \{0, 1\}$ determined by $z \in C_{i(z)}$. Hence, if one wants to prove the nonexistence of almost invariant, nontrivial colorings, it suffices to consider colorings where two colors are used.

If $S \subset G$ is a set of generators for G, a coloring is almost invariant if and only if for each $g \in S$ we have $c(x) = c(gx)$ for all but finitely many $x \in X$. This observation is of course particularly useful when G is finitely generated, which is the case when G is the mapping class group. Hence both G and X are countable. Also, it is easy to see that any almost invariant coloring of X can at most use finitely many colors: Assume without loss of generality that $c\colon X \to C$ is surjective, and for $z \in C$ let $X_z = c^{-1}(z)$; then $X = \bigsqcup_{z \in C} X_z$ is the partition of X associated to c. Next, choose some finite set of generators $g_1, \ldots, g_k$ of G. The almost invariance of the coloring implies that each g_i acts as a permutation of all but finitely many X_z, hence G acts as a permutation on all but finitely many of the subsets. If the partition consists of infinitely many subsets, this contradicts the assumption that G acts transitively on X.

Theorem 7.1 *Let D be a (nonempty) multicurve on Σ, and let $X = \Gamma D$ be the mapping class group orbit of D. There are no nontrivial almost invariant colorings of X.*

This is essentially the last ingredient in the proof of Theorem 1.1, and the proof will occupy the remaining part of this section. The assumption that Σ has genus at least 2 is essential. In Section 9, it is proved that when Σ is a closed torus and X is the set of all nonseparating simple closed curves, there exist almost invariant colorings using arbitrarily many colors.

Referring to [10], we see that Theorem 7.1 can also be formulated as follows: The pair (Γ, Γ_D) has only one end, where Γ_D is the stabilizer of D in Γ.

7.1 Interesting pairs

We will assume that the elements of X have been colored red and blue, and then prove that one of these colors has only been used a finite number of times. To this end, an *interesting pair* is a pair (τ_γ, D) where τ_γ is a Dehn twist in a curve γ and $D \in X$ is a multicurve such that $\tau_\gamma D \neq D$ (equivalently, $i(\gamma, D) > 0$). Since τ_γ changes the color of only finitely many multicurves, the multicurves $\tau_\gamma^n D$ all have the same color for all sufficiently large values of n. This color is called the *future* of the interesting pair (τ_γ, D), denoted $\mathrm{fut}(\tau_\gamma, D)$. Similarly, we may consider the *past* $\mathrm{pas}(\tau_\gamma, D)$ of an interesting pair; the common color of all multicurves $\tau_\gamma^{-n} D$ for sufficiently large n. We will also need to consider pairs of the form (τ_γ^{-1}, D); the same definition of future and past applies to these, and clearly $\mathrm{fut}(\tau_\gamma^{\pm 1}, D) = \mathrm{pas}(\tau_\gamma^{\mp 1}, D)$.

Lemma 7.2 *For any interesting pair (τ_α, D), we have*

$$\mathrm{pas}(\tau_\alpha^{-1}, D) = \mathrm{fut}(\tau_\alpha, D) = \mathrm{pas}(\tau_\alpha, D) = \mathrm{fut}(\tau_\alpha^{-1}, D) \tag{8}$$

Proof It suffices to prove the middle identity. We may find a nonseparating simple closed curve β different and disjoint from α such that (τ_β, D) is also interesting. To see this, let δ be a component of D for which $\tau_\alpha \delta \neq \delta$, and assume that α and δ are represented by geodesics with respect to some choice of hyperbolic metric. Cutting Σ along α then yields a (possibly nonconnected) surface with geodesic boundary, in which δ is a number of properly embedded hyperbolic arcs. At least one of the connected components of the cut surface has genus at least 1, so in this component we may find a closed geodesic β, not parallel to a boundary component, intersecting one of the δ–arcs. In the original surface, β is still a geodesic intersecting the geodesic δ; hence $\tau_\beta \delta \neq \delta$ and (τ_β, D) is interesting.

Next, since τ_α and τ_β commute, we see that $\tau_\alpha^n \tau_\beta^m D$ is an $\mathbf{Z} \times \mathbf{Z}$–indexed family of distinct multicurves. By assumption, both τ_α and τ_β change the color of finitely many multicurves. Hence, outside some bounded region in $\mathbf{Z} \times \mathbf{Z}$, moving from one multicurve to a neighbour does not change the color, and since we can connect the future of (τ_α, D) to its past using such moves, the claim follows. □

From now on, we will only consider the future.

Lemma 7.3 *Assume that α and β are simple closed curves with $i(\alpha,\beta)\le 1$, and that D is a multicurve such that (τ_α, D), (τ_β, D) are interesting pairs. Then* $\mathrm{fut}(\tau_\alpha, D) = \mathrm{fut}(\tau_\beta, D)$.

Proof If $i(\alpha,\beta)=0$ the result follows from the proof of Lemma 7.2.

Now assume $i(\alpha,\beta)=1$. Then $\alpha\cup\beta$ is contained in a subsurface Σ' of genus 1 with one boundary component γ. If D can not be isotoped to be contained entirely in Σ', either some component of D intersects γ essentially, or some component of D lives in the complement of Σ'. In the former case, it is clear that (τ_γ, D) is interesting, so the $i=0$ case implies $\mathrm{fut}(\tau_\alpha, D)=\mathrm{fut}(\tau_\gamma, D)=\mathrm{fut}(\tau_\beta, D)$. In the latter case, use the fact that the complement of Σ' has genus at least 1 to find a simple closed curve intersecting D essentially.

Otherwise, D lives entirely in Σ'. Let D_α denote any component of D on which τ_α acts nontrivially. Then D_α is a simple closed curve in a torus with one boundary component. Since D_α is not a parallel copy of the boundary component, it must be a nonseparating curve not parallel to α. Using (oriented versions of) α and β as a basis for $H_1(\Sigma')$, the coordinates of (an oriented version of) D_α must be a pair (p,q) with $(p,q)\neq(1,0)$. Clearly, any other component of D is forced to be either parallel to the boundary component of Σ' or to D_α. The only way that τ_β can act on some component of D is then that τ_β acts on D_α; hence also $(p,q)\neq(0,1)$.

Consider the schematic picture of Σ' on Figure 2, where the boundary component is

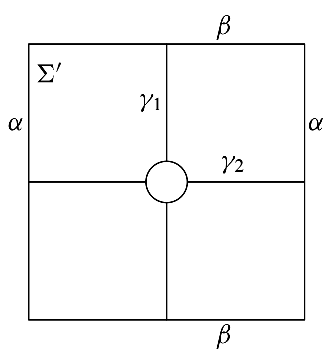

Figure 2: A torus with one boundary component

the circle in the center and α and β are the sides of the square.

We construct two disjoint simple closed curves γ_1, γ_2 as follows: Draw two essential, disjoint arcs in Σ' with the endpoints on the boundary component, and use the fact that the complement of Σ' has genus at least 1 to close them up in such a way that they are disjoint and not homotopic to a curve contained in Σ'. By the above description of D_α, (τ_{γ_n}, D) are both interesting pairs. Now the $i = 0$ case implies that

$$\mathrm{fut}(\tau_\alpha, D) = \mathrm{fut}(\tau_{\gamma_1}, D) = \mathrm{fut}(\tau_{\gamma_2}, D) = \mathrm{fut}(\tau_\beta, D). \qquad \square$$

The next proposition extends the above lemma to $i(\alpha, \beta) \leq 2$, but its proof is rather technical. Also, as explained in the comments following the proof, it is in fact not needed when one is only interested in surfaces with at most one boundary component.

Proposition 7.4 *Assume that α and β are simple closed curves with $i(\alpha, \beta) = 2$, and that D is a multicurve such that (τ_α, D) and (τ_β, D) are interesting. Then* $\mathrm{fut}(\tau_\alpha, D) = \mathrm{fut}(\tau_\beta, D)$.

Proof Let N be a regular neighbourhood of $\alpha \cup \beta$. We distinguish four cases.

(1) At least one of α and β is nonseparating in N.

(2) Both α and β are separating in N, but nonseparating in Σ.

(3) Both α and β are separating in N, but one is nonseparating in Σ.

(4) Both α and β are separating in Σ.

In case (1), assume without loss of generality that α is nonseparating. This means that when cutting N along α, there is at least one arc b of β connecting the two sides of α. Now construct two curves γ_1, γ_2 as follows: Make two parallel copies of b and close them up using arcs going in opposite directions along α. Applying small isotopies in a tubular neighbourhood of α we obtain a situation as depicted in Figure 3. We observe that each γ_n intersects α in exactly one point, and also they intersect each other in exactly one point p. Furthermore, since $i(\alpha, \beta) = 2$, the arc b does not start and end at the same point of α, so we have $i(\gamma_n, \beta) = 1$ for $n = 1, 2$.

Now let D_α be some component of D on which τ_α acts nontrivially. We claim that at least one of γ_1 and γ_2 intersects D_α essentially. Assume the contrary, and orient γ_1 and γ_2 oppositely along b. Choose geodesic representatives γ_1', γ_2' and D_α' of these three curves. Then γ_n' is disjoint from D_α', and necessarily γ_1' and γ_2' intersect transversally in a single point p'. But then $(\gamma_1'\gamma_2')_{p'} \in \pi_1(\Sigma, p')$ is a representative of the free homotopy class of (an oriented version of) α which does not intersect D_α', implying that $i(D_\alpha, \alpha) = 0$. This contradicts the choice of D_α.

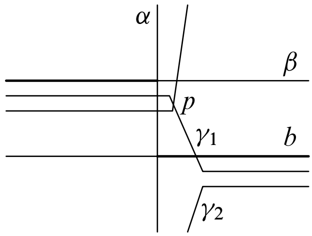

Figure 3: When α is nonseparating in N, the two sides of α are connected by an arc of β.

So one of the pairs (τ_{γ_n}, D) is interesting, and by Lemma 7.3 we have

$$\mathrm{fut}(\tau_\alpha, D) = \mathrm{fut}(\tau_{\gamma_n}, D) = \mathrm{fut}(\tau_\beta, D).$$

This ends case (1).

In cases (2)–(4), notice that N is necessarily a sphere with four holes, and α and β divide N into two pairs of pants in two different ways. Denote the boundary components of N by γ_i, $i = 0, 1, 2, 3$, such that γ_1, γ_2 are on one side of α and γ_0, γ_3 on the other, and such that γ_0, γ_1 are on one side of β and γ_2, γ_3 on the other. Schematically we have Figure 4(a).

Throughout the rest of the proof, we assume that α, β, γ_i, $i = 0, 1, 2, 3$, denote geodesic representatives for their isotopy classes. Also, we let δ be the geodesic representative of some component of D on which τ_α acts nontrivially. If δ does not live entirely in N, a twist in one of the boundary components acts nontrivially on δ, and since this boundary component is disjoint from α and β we are done by Lemma 7.3. Otherwise, δ is a separating curve in N which is not parallel to a boundary component. Clearly δ can not be parallel to β, since in that case D could not consist of any component on which τ_β acts nontrivially. Thus δ is different from both α and β.

In case (2), it is not hard to see that at least one of the "opposite" pairs γ_1, γ_3 and γ_0, γ_2 can be connected by an arc in the complement of N. Take two parallel copies of this arc, and close them up by arcs intersecting each other, α and β exactly once as in Figure 4(b) (the two connecting arcs are related by a twist in α). We may then argue exactly as in case (1) to see that the twist in at least one of these simple closed curves acts nontrivially on the multicurve in question.

In case (3), assume without loss of generality that β is separating and α is nonseparating. This means that it is impossible to connect any of γ_0 and γ_1 to any of γ_2 and γ_3 in the complement of N. But then, since α is nonseparating, one may connect either γ_0 to γ_1

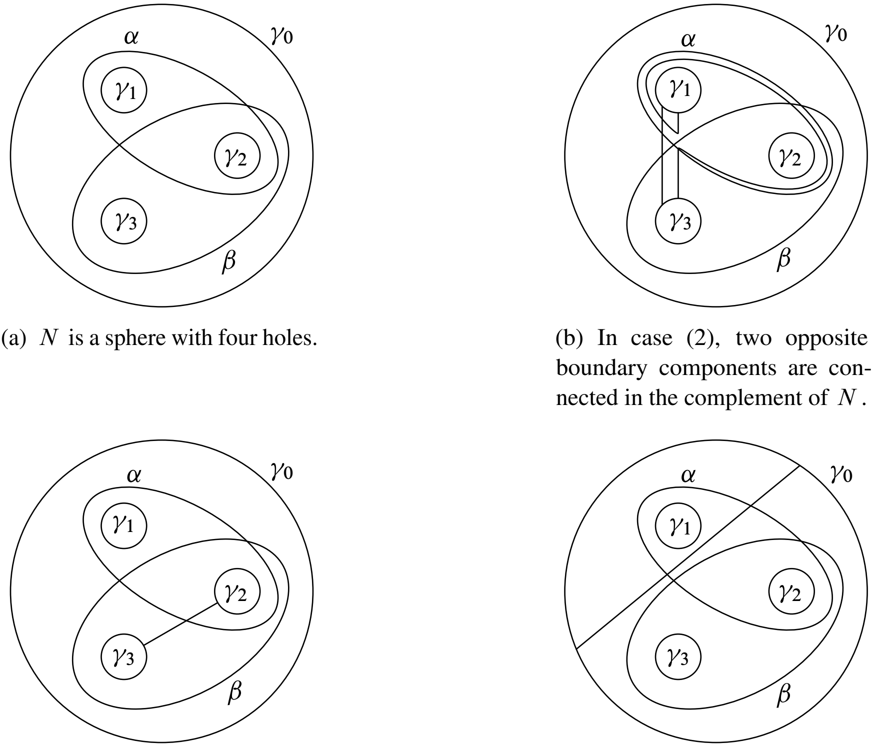

(a) N is a sphere with four holes.

(b) In case (2), two opposite boundary components are connected in the complement of N.

(c) In case (3), two "neighbouring" boundary components are connected in the complement of N.

(d) In case (4), there exists an essential arc in the complement of N starting and ending at the same boundary component.

Figure 4: There are four different topological cases when two curves intersect in two points.

or γ_2 to γ_3 in the complement of N. Assume without loss of generality that the latter is the case, and construct a simple closed curve γ disjoint from β intersecting γ_2, α and γ_3 exactly once each by composing the arc in the complement of N with an arc in N, as in Figure 4(c). Observe that the geodesic representative of γ necessarily intersects γ_2, α and γ_3 exactly once and is disjoint from β, so this representative contains a subarc in N starting at γ_2 and ending at γ_3. We now claim that this arc intersects δ (recall that δ has been chosen to be a geodesic). Assume the contrary. Then δ is a simple closed curve in the surface obtained by cutting N along this arc, which is a pair of pants. The "legs" are γ_0 and γ_1, whereas the "waist" is composed of four segments; two copies of the connecting arc and the remaining boundary components (cut open).

Since δ is simple, it is parallel to one of the boundary components of the pair of pants. But δ is certainly not parallel to any of the original boundary components, nor is it parallel to the "waist", since the latter is parallel to β. This contradiction implies that (τ_γ, D) is an interesting pair, and since γ is disjoint from β and intersects α in a single point, Lemma 7.3 yields the desired result,

$$\mathrm{fut}(\tau_\alpha, D) = \mathrm{fut}(\tau_\gamma, D) = \mathrm{fut}(\tau_\beta, D).$$

Finally, in case (4), none of the four boundary components of N can be connected in the complement of N. This means that at least one of the connected components of $\Sigma - N$ must have positive genus. Assume without loss of generality that the component Σ_0 bounded by γ_0 has positive genus. Now take some nonseparating, essential arc in Σ_0 with its endpoints on γ_0 and compose it with some essential arc in N disjoint from β and intersecting α in exactly two points (cf Figure 4(d)) to obtain a nonseparating curve γ in Σ. We claim that τ_γ acts nontrivially on δ, ie that the arc in N intersects δ essentially. To see this, we argue as in case (3) above. Observe that γ has geometric intersection number 2 with α and γ_0. Hence, the geodesic representative of γ intersects α and γ_0 exactly twice, so this geodesic contains a subarc in N looking as the one depicted in Figure 4(d). We claim that this arc intersects δ. If this were not the case, we may cut N along this arc to obtain a cylinder (bounded by one of the original boundary components and a curve coming from the cut) and a pair of pants (bounded by two of the original boundary components and a curve from the cut), and δ lives completely in one of these. Since δ is not parallel to any of the boundary components of N, we conclude that δ is parallel to the third boundary component of the pair of pants. But this third boundary component is clearly parallel to β, which contradicts the fact that D does not contain any component parallel to β. Hence (τ_γ, D) is interesting, and since γ is nonseparating and intersects α in two points, by case (3) and Lemma 7.3 we have

$$\mathrm{fut}(\tau_\alpha, D) = \mathrm{fut}(\tau_\gamma, D) = \mathrm{fut}(\tau_\beta, D),$$

which finishes the last case. □

Now we turn to the (finite) presentation of the mapping class group given by Gervais [7], where the generators are twists in certain curves. A key property of this presentation is that any two curves involved intersect each other in at most two points. It should be pointed out, however, that if one is only interested in surfaces with at most one boundary component, a much earlier result by Wajnryb [12] yields a presentation where each pair of curves intersect in at most one point. In this case, one does not need the rather technical Proposition 7.4 above in the following (simply replace all references to [7] by [12] and all occurences of "at most two points" by "at most one point").

Proposition 7.5 *Let S denote the set of curves from [7] such that $\{\tau_\sigma \mid \sigma \in S\}$ generate Γ. Let $\alpha, \beta \in S$ be two of these curves, and let $D_1, D_2 \in X$ be multicurves such that (τ_α, D_1) and (τ_β, D_2) are interesting. Then*

$$\operatorname{fut}(\tau_\alpha, D_1) = \operatorname{fut}(\tau_\beta, D_2).$$

Proof We may find a sequence of curves $\eta_1, \eta_2, \ldots, \eta_n \in S$ and exponents $\varepsilon_i = \pm 1$ such that, writing $\tau_i = \tau_{\eta_i}^{\varepsilon_i}$, $\tau_n \cdots \tau_2 \tau_1 D_1 = D_2$. For each $1 \leq i \leq n$ we may assume that $(\tau_i, \tau_{i-1} \cdots \tau_1 D_1)$ is interesting; otherwise we may simply omit the corresponding τ_i. Now using alternately the fact that η_i and η_{i+1} intersect in at most two points and the obvious fact that $\operatorname{fut}(\tau_\gamma, D) = \operatorname{fut}(\tau_\gamma, \tau_\gamma D)$ for any interesting pair (τ_γ, D), we obtain a sequence of identities

$$\begin{aligned}
\operatorname{fut}(\tau_1, D_1) &= \operatorname{fut}(\tau_1, \tau_1 D_1) = \operatorname{fut}(\tau_2, \tau_1 D_1) \\
&= \operatorname{fut}(\tau_2, \tau_2 \tau_1 D_1) = \operatorname{fut}(\tau_3, \tau_2 \tau_1 D_1) \\
&\ \vdots \\
&= \operatorname{fut}(\tau_{n-1}, \tau_{n-1} \cdots \tau_2 \tau_1 D_1) = \operatorname{fut}(\tau_n, \tau_{n-1} \cdots \tau_2 \tau_1 D_1) \\
&= \operatorname{fut}(\tau_n, \tau_n \cdots \tau_2 \tau_1 D_1) = \operatorname{fut}(\tau_n, D_2)
\end{aligned}$$

which may be augmented by the identities $\operatorname{fut}(\tau_\alpha, D_1) = \operatorname{fut}(\tau_1, D_1)$ and $\operatorname{fut}(\tau_n, D_2) = \operatorname{fut}(\tau_\beta, D_2)$ to obtain the desired result. □

Lemma 7.6 *Let $f \in \Gamma$ be any diffeomorphism, and (τ_α, D) an interesting pair. Then $(\tau_{f(\alpha)}, fD)$ is also interesting and $\operatorname{fut}(\tau_\alpha, D) = \operatorname{fut}(\tau_{f(\alpha)}, fD)$.*

Proof Recall that $f \circ \tau_\alpha \circ f^{-1} = \tau_{f(\alpha)}$. Hence $\tau_{f(\alpha)}(fD) = f(\tau_\alpha D) \neq fD$, so $(\tau_{f(\alpha)}, fD)$ is interesting. Also we have

$$\tau_{f(\alpha)}^n = f \circ \tau_\alpha^n \circ f^{-1},$$

so

$$\tau_{f(\alpha)}^n(fD) = f(\tau_\alpha^n D).$$

Since the different multicurves $\tau_\alpha^n D$ have the same color for all sufficiently large n, and since f changes the color of only finitely many multicurves, the result follows. □

Proposition 7.7 *All interesting pairs (τ_γ, D) where γ is a nonseparating curve have the same future.*

Proof Let τ_α be a twist on a nonseparating curve which is part of the generating set for Γ from [7]. Then Proposition 7.5, with $\alpha = \beta$, implies that the future is a property of τ_α alone, and not of the particular multicurve on which τ_α acts. If γ is any nonseparating curve, choose a diffeomorphism of Σ carrying γ to α and apply Lemma 7.6. □

Now we are ready to prove the nonexistence of almost invariant colorings.

Proof of Theorem 7.1 Choose a finite set $\alpha_1, \dots, \alpha_N$ of nonseparating curves such that the twists in these curves generate Γ (we do not require that these intersect pairwise in at most two points). To be concrete, assume that the common future (cf Proposition 7.7) of all interesting pairs (τ_γ, D) with γ nonseparating is red. We must then prove that only finitely many multicurves are blue. Let $B \subset X$ be the set of blue multicurves. For each blue multicurve $D \in B$, choose a generator τ_{α_k} such that (τ_{α_k}, D) is interesting (this must be possible since the action is transitive and the τ_{α_k} generate Γ). This defines a map $f \colon B \to \{1, 2, \dots, N\}$. We claim that for each $k \in \{1, \dots, N\}$, the preimage $f^{-1}(k)$ is finite.

To see this, for each $D \in f^{-1}(k)$ consider the "τ_{α_k}–string through D", ie the set $s_k(D) = \{\tau_{\alpha_k}^n D \mid n \in \mathbf{Z}\}$. Let B_k be the union of the blue multicurves occurring in these strings, ie

$$B_k = \bigcup_{D \in f^{-1}(k)} (s_k(D) \cap B),$$

so that $f^{-1}(k) \subseteq B_k$. There are only finitely many blue multicurves in each string by Proposition 7.7 and Lemma 7.2. Since τ_{α_k} changes the color of at least one multicurve in each string (since the strings contain both blue and red multicurves), there can be only finitely many strings by the almost invariance of the coloring. Hence, there are only finitely many blue multicurves. □

8 Proof of the main theorem

Now we have all the tools we need.

Proof of Theorem 1.1 By the isomorphism (1) and the splitting (2), it suffices to prove that each summand $H^1(\Gamma, M_D)$ vanishes. By Proposition 6.1, we need only show that the map ι_* is injective. By the exact sequence (4), this is equivalent to proving that $H^0(\Gamma, \widehat{M}_D) \to H^0(\Gamma, \widehat{M}_D / M_D)$ is surjective.

Now, an invariant element of $\widehat{M}_D / M_D$ is represented by an element $v \in \widehat{M}_D = \mathrm{Map}(\Gamma D, \mathbf{C})$ such that for each $g \in \Gamma$ we have $v - gv \in M_D$. Since $(v - gv)(E) = v(E) - v(g^{-1}E)$ for $E \in \Gamma D$, we see that this must be zero for all but finitely many

$E \in \Gamma D$. In other words, v must be an almost invariant $\mathbf{C}$–coloring of ΓD in the above language, and since by Theorem 7.1 no nontrivial almost invariant colorings of ΓD exist, we conclude that v is almost constant, ie all but finitely many elements of ΓD is mapped to the same complex number z. But then v represents the same element of $\widehat{M}_D/M_D$ as the constant linear combination $\sum_{E \in \Gamma D} zE$, and hence $H^0(\Gamma, \widehat{M}_D) \to H^0(\Gamma, \widehat{M}_D/M_D)$ is in fact surjective. □

9 The genus one case

When Σ is a closed torus, it is well-known that $\Gamma \cong \mathrm{SL}_2(\mathbf{Z})$. A multicurve necessarily consists of some number of parallel copies of the same nonseparating simple closed curve. Hence, a reduced multicurve is simply a nonseparating simple closed curve, and $H^1(\Gamma, \mathcal{O}(\mathcal{M}))$ is a countable direct sum

$$H^1(\Gamma, \mathcal{O}(\mathcal{M})) \cong \bigoplus_{n \in \mathbf{Z}_+} H^1(\Gamma, M_{D_n}), \tag{9}$$

where each summand is isomorphic to $H^1(\Gamma, M_\gamma)$, where γ is some fixed nonseparating curve.

Let $X = \Gamma\gamma$ be the mapping class group orbit of γ. We may identify X with the set of unoriented torus knots, ie the set P of pairs (p,q), $p,q \in \mathbf{Z}$ and $\gcd(p,q) = 1$, where we identify the pairs (p,q) and $(-p,-q)$ (since the curves are not oriented). The action of the mapping class group is then simply given by the usual action of $\mathrm{SL}_2(\mathbf{Z})$ on pairs of relatively prime integers, and the central element $-I$ acts trivially, so we are really dealing with an action of $\mathrm{PSL}_2(\mathbf{Z})$.

As generators for $\mathrm{SL}_2(\mathbf{Z})$ we choose

$$S = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \quad \text{and} \quad R = \begin{bmatrix} 0 & -1 \\ 1 & 1 \end{bmatrix}.$$

Then $S^2 = R^3 = I$ in $\mathrm{PSL}_2(\mathbf{Z})$. Letting

$$\begin{aligned} X_1 &= \{(p,q) \mid p \geq 1, q \geq 0\} \\ X_2 &= \{(p,q) \mid q > -p \geq 0\} \\ X_3 &= \{(p,q) \mid -p \geq q > 0\} \end{aligned}$$

it is easy to see that $X_1 \cup X_2 \cup X_3 = X$, and one also verifies that $SX_1 = X_2 \cup X_3$, $RX_1 = X_2$, $RX_2 = X_3$.

Proposition 9.1 *Any point $(p,q) \in X$ with $p,q > 0$, can be reached from $(1,1)$ by applying a unique sequence of elements of $\mathrm{SL}_2(\mathbf{Z})$ of the form $S^{-1}R^k$, where k is 1 or 2.*

Proof For existence, we will use induction on $\max(p,q)$. For $\max(p,q) = 1$ we have $p = q = 1$, in which case the claim is obvious (choose the empty sequence). If $\max(p,q) > 1$, p and q are different since $\gcd(p,q) = 1$. If $p > q$, put

$$(p',q') = R^{-1}S(p,q) = R^{-1}(-q,p) = (p-q,q)$$

while if $q > p$, put

$$(p',q') = R^{-2}S(p,q) = R^{-2}(-q,p) = (p,q-p).$$

In both cases, clearly $1 \le p',q'$ and $\max(p',q') < \max(p,q)$, so there exists $\gamma' = S^{-1}R^{k_{n-1}} \cdots S^{-1}R^{k_1}$ with $\gamma'(1,1) = (p',q')$. Then $\gamma = S^{-1}R^{k_n}\gamma'$ where $k_n = 1$ if $p > q$ and $k_n = 2$ if $p < q$ is an element of $\mathrm{PSL}_2(\mathbf{Z})$ of the desired form.

To prove uniqueness, choose (p,q) with $\max(p,q)$ minimal such that there are two different strings

$$\gamma_1 = S^{-1}R^{k_n}S^{-1}R^{k_{n-1}} \cdots S^{-1}R^{k_1}$$
$$\gamma_2 = S^{-1}R^{\ell_m}S^{-1}R^{\ell_{m-1}} \cdots S^{-1}R^{\ell_1}$$

satisfying $\gamma_i(1,1) = (p,q)$. Then $S(p,q)$ is a point in $X_2 \cup X_3$ which is obtained by applying R^{k_n} to some point of X_1 and also by applying R^{ℓ_m} to some (possibly other) point of X_1. But since $S(p,q)$ is an element of exactly one of $X_2 = RX_1$ and $X_3 = R^2X_1$, this implies that $k_n = \ell_m$. Continuing this way, we only need to show that there is no nontrivial string

$$\gamma = S^{-1}R^{k_n}S^{-1}R^{k_{n-1}} \cdots S^{-1}R^{k_1}$$

such that $\gamma(1,1) = (1,1)$. But this is trivial by observing that each element of the form $S^{-1}R^k$ strictly increases the max–norm of any point (p,q) with $p,q \ge 1$ (since $S^{-1}R(p,q) = (p+q,q)$ and $S^{-1}R^2(p,q) = (p,p+q)$). □

This proposition allows us to label the vertices of an infinite binary tree T as follows. The root is labelled by $(1,1)$, and all remaining vertices are labelled according to the rule: If a vertex v is reached by going "left" from the immediate predecessor, the label of v is obtained by applying $S^{-1}R$ to the label of its predecessor; otherwise the label is obtained by applying $S^{-1}R^2$ (see Figure 5).

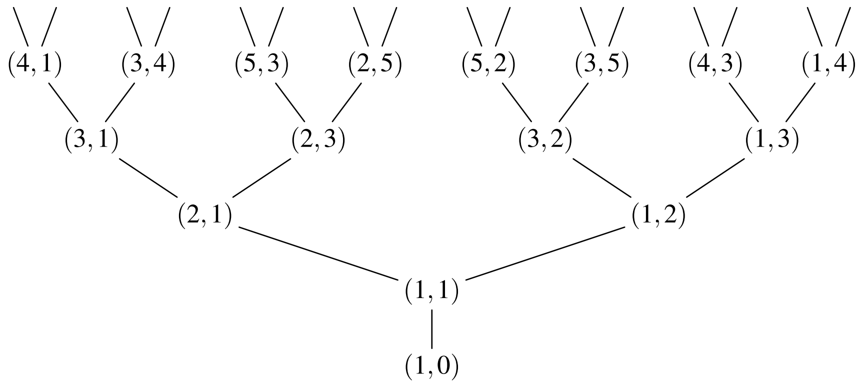

Figure 5: An infinite binary tree labelled by the points of X_1

Now add a single vertex below the root and label this by $(1,0)$. This gives, by Proposition 9.1, a 1–1 correspondence between the vertices of T and the points in X_1, and from now on we shall refer to a vertex and its label interchangeably.

By the level of a vertex of T we mean its distance from $(1,1)$ (the level of $(1,0)$ may be taken to be -1); there are 2^k vertices at level k for each $k \geq 0$, and also exactly 2^k vertices at level $< k$.

Proposition 9.2 *For each $k \geq 0$, there is an almost invariant coloring of X using 2^k different colors.*

Proof We start by coloring the subset X_1 by coloring the vertices of the tree. Assign different colors to the 2^k vertices at level k, and for each of these vertices assign the same color to all descendants. The remaining 2^k points of X_1 may be colored arbitrarily.

To obtain a coloring of all of X, we insist that the coloring is completely invariant under R. This gives a well-defined coloring, since X_1 is a complete set of representatives of the R–orbits of X. In order to see that this coloring is almost invariant under $\mathrm{PSL}_2(\mathbf{Z})$, it suffices to check that the other generator S changes the color of only finitely many points of X. Since S has order two in $\mathrm{PSL}_2(\mathbf{Z})$, S changes the color of p if and only if it changes the color of Sp. Hence we need only check that S changes the color of finitely many elements of X_1. But for any vertex v of T of level $k+1$ or higher, applying S to the label of v yields by construction a point of X_2 or X_3 which has the same color as the predecessor of v; hence S does not change the color of labels placed at level $k+1$ or higher, and thus S changes the color of at most $2 \cdot 2^{k+1}$ points of X. □

We are also able to give a classification of all almost invariant colorings.

Proposition 9.3 *Any almost invariant coloring of X is equivalent to (a coloring which is a simplification of) a coloring of the form constructed in Proposition 9.2.*

Proof Let $c\colon X \to C$ be some almost invariant coloring of X. Since R changes the color of only finitely many points of X, it changes the color of only finitely many points $x_1, \dots, x_N$ of X_1. Now we change c into an equivalent coloring c' by putting $c'(Rx_i) = c'(R^2 x_i) = c(x_i)$, and $c' = c$ otherwise. Then c' is by construction completely invariant under R. Now since c' is almost invariant, there are only finitely many points of X_1 whose c'–color changes under S. Choose K such that the color of any label placed at level $k > K$ is unchanged under S. This, together with the R–invariance of c', implies that each label at level K has the same color as any of its descendants, and hence c' is (a simplification of) a coloring using 2^K different colors. □

Theorem 9.4 *The cohomology group $H^1(\Gamma, \mathcal{O}(\mathcal{M}))$ is infinite-dimensional.*

Proof By the isomorphism (9), it suffices to prove that $H^1(\Gamma, M_\gamma)$ is nonzero. By Proposition 9.2, the space $H^0(\Gamma, \widehat{M}_\gamma / M_\gamma)$ is infinite-dimensional, and since $H^0(\Gamma, \widehat{M}_\gamma) \cong \mathbf{C}$ is 1–dimensional, the exact sequence (4) implies that $H^1(\Gamma, M_\gamma)$ is in fact infinite-dimensional. □

It is not hard to give an explicit example of a nonzero cohomology class $[u] \in H^1(\Gamma, \mathcal{O}(\mathcal{M}))$. Consider two simple closed curves α, β intersecting transversely in a single point. Then $\tau_\alpha \tau_\beta$ corresponds to R above, and $\tau_\alpha \tau_\beta \tau_\alpha$ corresponds to S. We put

$$u(\tau_\alpha) = \alpha - \beta$$
$$u(\tau_\beta) = \beta - \alpha,$$

and by expanding using the cocycle condition, one checks that $u(\tau_\alpha \tau_\beta \tau_\alpha) = u(\tau_\beta \tau_\alpha \tau_\beta)$ and $u((\tau_\alpha \tau_\beta \tau_\alpha)^4) = 0$. This means that u extends to a well-defined cocycle $u\colon \Gamma \to M_\gamma$. Clearly, for any element $m \in M_\gamma$, the coefficient of α in $\delta m(\tau_\alpha) = (1 - \tau_\alpha) m$ is 0, so u defines a nonzero element of $H^1(\Gamma, M_\gamma)$.

References

[1] **J E Andersen**, *Hitchin's connection, Toeplitz operators and symmetry invariant deformation quantization* arXiv:math.DG/0611126

[2] **J E Andersen**, **N L Gammelgaard**, *Hitchin's projectively flat connection, Toeplitz operators and the asymptotic expansion of TQFT curve operators* arXiv:0903.4091

[3] **J E Andersen**, **R Villemoes**, *Degree one cohomology with twisted coefficients of the mapping class group* arXiv:0710.2203

[4] **K S Brown**, *Cohomology of groups*, Graduate Texts in Math. 87, Springer, New York (1982) MR672956

[5] **D Bullock**, **C Frohman**, **J Kania-Bartoszyńska**, *Understanding the Kauffman bracket skein module*, J. Knot Theory Ramifications 8 (1999) 265–277 MR1691437

[6] **L Charles**, **J Marche**, *Multicurves and regular functions on the representation variety of a surface in* $\mathrm{SU}(2)$ arXiv:0901.3064

[7] **S Gervais**, *A finite presentation of the mapping class group of a punctured surface*, Topology 40 (2001) 703–725 MR1851559

[8] **W M Goldman**, *The complex-symplectic geometry of* $\mathrm{SL}(2,\mathbb{C})$*–characters over surfaces*, from: "Algebraic groups and arithmetic", (S G Dani, G Prasad, editors), Tata Inst. Fund. Res., Mumbai (2004) 375–407 MR2094117

[9] **R C Penner**, **J L Harer**, *Combinatorics of train tracks*, Annals of Math. Studies 125, Princeton Univ. Press (1992) MR1144770

[10] **P Scott**, *Ends of pairs of groups*, J. Pure Appl. Algebra 11 (1977) 179–198 MR487104

[11] **A R Skovborg**, *The moduli space of flat connections on a surface – Poisson structures and quantization*, PhD thesis, University of Aarhus (2006) Available at http://www.imf.au.dk/publs?id=623

[12] **B Wajnryb**, *A simple presentation for the mapping class group of an orientable surface*, Israel J. Math. 45 (1983) 157–174 MR719117

Department of Mathematics, University of Aarhus
8000 Aarhus C, Denmark

andersen@imf.au.dk, math@rasmusvillemoes.dk

http://home.imf.au.dk/andersen/, http://www.rasmusvillemoes.dk

Received: 24 October 2008 Revised: 13 May 2009