

THE DIMENSION OF THE TORELLI GROUP

MLADEN BESTVINA, KAI-UWE BUX, AND DAN MARGALIT

1. INTRODUCTION

Let S be a surface (unless specified otherwise, we take all surfaces to be connected, orientable, and of finite type). Let $\text{Mod}(S)$ be the *mapping class group* of S , defined as $\pi_0(\text{Homeo}^+(S))$, where $\text{Homeo}^+(S)$ is the group of orientation-preserving homeomorphisms of S . The *Torelli group* $\mathcal{I}(S)$ is the kernel of the natural action of $\text{Mod}(S)$ on $H_1(S, \mathbb{Z})$. When S is S_g , the closed surface of genus g , this action is symplectic (it preserves the algebraic intersection number), and it is a classical fact that $\text{Mod}(S_g)$ surjects onto $\text{Sp}(2g, \mathbb{Z})$. All of this information is encoded in the following short exact sequence:

$$1 \rightarrow \mathcal{I}(S_g) \rightarrow \text{Mod}(S_g) \rightarrow \text{Sp}(2g, \mathbb{Z}) \rightarrow 1.$$

Cohomological dimension. For a group G , we denote by $\text{cd}(G)$ its *cohomological dimension*, which is the supremum over all n so that there exists a G -module M with $H^n(G, M) \neq 0$.

Theorem A. *For $g \geq 2$, we have $\text{cd}(\mathcal{I}(S_g)) = 3g - 5$.*

Since $\text{Mod}(S_0) = 1$, we have $\mathcal{I}(S_0) = 1$. Also, it is a classical fact that $\text{Mod}(S_1) \cong \text{Sp}(2, \mathbb{Z}) = \text{SL}(2, \mathbb{Z})$, and so $\mathcal{I}(S_1)$ is also trivial.

The lower bound for Theorem A was already given by Mess [36] in an unpublished paper from 1990.

Theorem 1.1 (Mess). *For $g \geq 2$, we have $\text{cd}(\mathcal{I}(S_g)) \geq 3g - 5$.*

Theorem 1.1 is proven by constructing a subgroup of $\mathcal{I}(S_g)$ that is a Poincaré duality subgroup of dimension $3g - 5$. We recall Mess's proof in Section 9. Mess [36] explicitly asked if the statement of Theorem A is true.

We also study the *Johnson kernel*, denoted $\mathcal{K}(S_g)$, which is the subgroup of $\mathcal{I}(S_g)$ generated by Dehn twists about separating curves. Birman [5] showed that $\mathcal{K}(S_2)$ is equal to $\mathcal{I}(S_2)$. On the other hand, Johnson [27] proved that the index of $\mathcal{K}(S_g)$ in $\mathcal{I}(S_g)$ is infinite when $g \geq 3$, answering a question of Birman.

The following theorem answers a question of Farb [15, Problem 5.9].

Theorem B. *For $g \geq 2$, we have $\text{cd}(\mathcal{K}(S_g)) = 2g - 3$.*

Received by the editors September 7, 2007.

2000 *Mathematics Subject Classification.* Primary 20F34; Secondary 57M07.

Key words and phrases. Mapping class group, Torelli group, Johnson kernel, cohomological dimension, complex of minimizing cycles.

The first and third authors gratefully acknowledge support by the National Science Foundation.

©2009 American Mathematical Society
Reverts to public domain 28 years from publication

$\mathcal{K}(S_g)$ has a free abelian subgroup of rank $2g - 3$ when $g \geq 2$ (see Section 9). From the inclusions $\mathbb{Z}^{2g-3} < \mathcal{K}(S_g) < \mathcal{I}(S_g)$, and Theorem A, it immediately follows that $2g - 3 \leq \text{cd}(\mathcal{K}(S_g)) \leq 3g - 5$ (see Fact 6.5 below).

We can think of Theorems A and B as giving the smallest possible dimensions of Eilenberg–MacLane spaces for $\mathcal{I}(S_g)$ and $\mathcal{K}(S_g)$, the so-called *geometric dimension*. Indeed, it is a theorem of Eilenberg–Ganea, Stallings, and Swan that if $\text{cd}(G) \neq 2$ for some group G , then $\text{cd}(G)$ is the same as the geometric dimension [14, 43, 44].

For a group with torsion, such as $\text{Mod}(S_g)$, the cohomological dimension is infinite. However, if a group G has a torsion free subgroup H of finite index, then we can define $\text{vcd}(G)$, the *virtual cohomological dimension* of G , to be $\text{cd}(H)$. It is a theorem of Serre [41, Théorème 1] that $\text{vcd}(G)$ does not depend on the choice of H . In 1986, Harer [19] proved that $\text{vcd}(\text{Mod}(S_g)) = 4g - 5$, so we see that there is a gap of g between $\text{cd}(\mathcal{I}(S_g))$ and $\text{vcd}(\text{Mod}(S_g))$.

Let $\mathcal{I}^k(S_g)$ be the subgroup of $\text{Mod}(S_g)$ consisting of elements that act trivially on a fixed $2k$ -dimensional symplectic subspace of $H_1(S, \mathbb{Z})$.

Conjecture 1.2. *For $g \geq 2$ and $0 \leq k \leq g$, $\text{vcd}(\mathcal{I}^k(S_g)) = 4g - 5 - k$.*

We discuss this conjecture further in Section 9.

The groups $\text{Mod}(S_g)$, $\mathcal{I}(S_g)$, and $\mathcal{K}(S_g)$ are the first three groups in a series of groups associated to S_g . Let $\Gamma = \pi_1(S_g)$, and let Γ_k be the k^{th} term of the lower central series of Γ , that is, $\Gamma_1 = \Gamma$ and $\Gamma_{k+1} = [\Gamma, \Gamma_k]$. The *Johnson filtration* of $\text{Mod}(S)$ is the sequence of groups $\{\mathcal{N}_k(S_g)\}$ given by

$$\mathcal{N}_k(S_g) = \ker\{\text{Mod}(S_g) \rightarrow \text{Out}(\Gamma/\Gamma_{k+1})\}.$$

It follows from the definitions that $\mathcal{N}_0(S_g) = \text{Mod}(S_g)$ and $\mathcal{N}_1(S_g) = \mathcal{I}(S_g)$, and it is a deep theorem of Johnson [30] that $\mathcal{N}_2(S_g) = \mathcal{K}(S_g)$.¹

For $g \geq 2$ and any $k \geq 0$, Farb [15] gives the lower bound $g - 1$ for $\text{cd}(\mathcal{N}_k(S_g))$ by constructing a free abelian subgroup of $\mathcal{N}_k(S_g)$ of rank $g - 1$. From this, and Theorem B, it follows that, for $g \geq 2$ and any $k \geq 2$, we have $g - 1 \leq \text{cd}(\mathcal{N}_k(S_g)) \leq 2g - 3$ (again apply Fact 6.5). The following question was already asked by Farb [15, Problem 5.9].

Question 1.3. What is $\text{cd}(\mathcal{N}_k(S_g))$ for $g \geq 3$ and $k \geq 3$?

Infinite generation of top homology. Besides the question of cohomological dimension, one would also like to know in which dimensions the homology is (in)finitely generated. For instance, if the second homology of a group is infinitely generated, then that group is not finitely presented. As of this writing, it is a central open problem to determine whether or not $\mathcal{I}(S_g)$ is finitely presented. We have the following theorem.

Theorem C. *For $g \geq 2$, the group $H_{3g-5}(\mathcal{I}(S_g), \mathbb{Z})$ is infinitely generated.*

In Kirby’s problem list, Problem 2.9(B) (attributed to Mess) is to find the largest $k = k_g$ so that $\mathcal{I}(S_g)$ has an Eilenberg–MacLane space with finite k -skeleton [1]. An immediate consequence of Theorem C is that $k_g \leq 3g - 6$. It follows from a theorem of Mess (see below) that $k_3 \leq 2$.

Theorem C also solves a problem of Farb [15, Problem 5.14].

We do not know if the top dimensional homology of $\mathcal{K}(S_g)$ is finitely generated for $g \geq 3$.

¹Johnson proved that $\mathcal{K}(S_g)$ is equal to the kernel of the Johnson homomorphism; this is different from saying that $\mathcal{K}(S_g) = \mathcal{N}_2(S_g)$, but the latter can be derived from the former [12].

Genus 2. It is a theorem of Stallings and Swan that a group of cohomological dimension 1 is free [43, 44]. Thus, Theorem A implies that $\mathcal{I}(S_2) = \mathcal{K}(S_2)$ is a free group. A celebrated theorem of Mess, Theorem D below, gives a much more precise picture of this group.

For the statement, note that each separating simple closed curve in S_2 gives rise to an algebraic splitting of $H_1(S_2, \mathbb{Z})$: the two subspaces are the ones spanned by the curves on the two different sides of the separating curve. Each component of the splitting is a 2-dimensional symplectic subspace of $H_1(S_2, \mathbb{Z})$, and so we call such a splitting a *symplectic splitting* of $H_1(S_2, \mathbb{Z})$. Note that two separating curves give rise to the same symplectic splitting if they differ by an element of $\mathcal{I}(S_2)$.

Theorem D (Mess). *$\mathcal{I}(S_2)$ is an infinitely generated free group, with one Dehn twist generator for each symplectic splitting of $H_1(S_2, \mathbb{Z})$.*

We give a new proof of this theorem in Section 7. While Mess's original proof is rooted in algebraic geometry, our proof is confined entirely to the realm of geometric group theory.

Strategy. All of our theorems cited above are proven by studying the actions of $\mathcal{I}(S_g)$ and $\mathcal{K}(S_g)$ on a new complex $\mathcal{B}(S_g)$, called the “complex of minimizing cycles”, which we construct in Section 2. Vertices of the complex correspond to integral representatives in S_g of some fixed element of $H_1(S_g, \mathbb{Z})$. The key feature is the following fact.

Theorem E. *The complex $\mathcal{B}(S_g)$ is contractible for $g \geq 1$.*

We give two proofs of Theorem E. In the first (in Sections 3 and 4), we use a nerve-cover argument to show that $\mathcal{B}(S_g)$ is homotopy equivalent to Teichmüller space. The second proof (in Section 5), which was discovered after the first version of this paper was distributed, uses a surgery argument to give an explicit contraction. Allen Hatcher has used similar ideas to show that a variant of our complex, the “cyclic cycle complex,” is contractible [21].

In light of Theorem 1.1 (and the discussion after Theorem B), Theorems A and B reduce to the following:

$$\text{cd}(\mathcal{I}(S_g)) \leq 3g - 5 \quad \text{and} \quad \text{cd}(\mathcal{K}(S_g)) \leq 2g - 3.$$

For a group G acting on a contractible cell complex X , Quillen [41, Proposition 11] proved that

$$\text{cd}(G) \leq \sup\{\text{cd}(\text{Stab}_G(\sigma)) + \dim(\sigma)\},$$

where σ ranges over the cells of X . We take X to be $\mathcal{B}(S_g)$, and we take G to be either $\mathcal{I}(S_g)$ or $\mathcal{K}(S_g)$. We then show in Section 6 that the right hand side is bounded above by $3g - 5$ for $\mathcal{I}(S_g)$ and by $2g - 3$ for $\mathcal{K}(S_g)$.

In genus 2, the picture is very explicit; we give a diagram of $\mathcal{B}(S_2)/\mathcal{I}(S_2)$ in Figure 8. Thus, we are able to deduce Theorem D by a careful analysis of this quotient. For a free group, infinite generation of the group is the same as infinite generation of its first homology. As such, Theorem D gives the genus 2 case of Theorem C.

For the proof of Theorem C we again appeal to the action of $\mathcal{I}(S_g)$ on $\mathcal{B}(S_g)$. It is a general fact that if a group G acts on a contractible cell complex X and satisfies the condition

$$\sup\{\text{cd}(\text{Stab}_G(\sigma)) + \dim(\sigma)\} = d$$

as above, then the d^{th} integral homology of the stabilizer in G of a vertex of X injects into $H_d(G, \mathbb{Z})$ (see Fact 8.2 in Section 8.1). In Section 8, we use an inductive argument (with base case Theorem D) to prove that the $(3g - 5)^{\text{th}}$ homology of the stabilizer in $\mathcal{I}(S_g)$ of a particular kind of vertex of $\mathcal{B}(S_g)$ is infinitely generated.

At the end of the paper we include an appendix that gives a classification of the multitwists that lie in the various terms of the Johnson filtration. This is needed for our proof of Theorem B.

History. The classical fact that $\mathcal{I}(S_1)$ is trivial was proven by Dehn [13] in the 1920s. In 1983, Johnson [29] proved that $\mathcal{I}(S_g)$ is finitely generated for $g \geq 3$, and in 1986, McCullough–Miller [34] proved that $\mathcal{I}(S_2)$ is not finitely generated. Mess [37] improved on this in 1992 by proving Theorem D. At the same time, Mess [37] showed that $H_3(\mathcal{I}(S_3), \mathbb{Z})$ is not finitely generated (Mess credits the argument to Johnson–Millson). In 2001, Akita [2] proved that $H_*(\mathcal{I}(S_g), \mathbb{Z})$ is not finitely generated for $g \geq 7$. Hain [17, Corollary 18] proved in 2002 that $H_4(\mathcal{I}(S_3), \mathbb{Z})$ is infinitely generated and also gave a new proof that $H_3(\mathcal{I}(S_3), \mathbb{Z})$ is infinitely generated.

Mess [36] proved Theorem 1.1, the lower bound for $\text{cd}(\mathcal{I}(S_g))$, in 1990. The authors [3] recently proved that $\text{cd}(\mathcal{I}(S_g)) < 4g - 5 = \text{vcd}(\text{Mod}(S_g))$. (Harer’s result already implied that $\text{cd}(\mathcal{I}(S_g)) \leq 4g - 5$.)

In the case of the Johnson kernel, less is known. As $\mathcal{K}(S_2) = \mathcal{I}(S_2)$, we have the theorem of Mess in this case. It is not known if $\mathcal{K}(S_g)$ or $H_1(\mathcal{K}(S_g))$ is finitely generated for $g \geq 3$.

In an earlier paper, the authors [3] studied the Torelli subgroup of $\text{Out}(F_n)$, that is, the group \mathcal{T}_n of outer automorphisms of a free group of rank n that act trivially on the first homology of the free group. The main results of that paper (that $\text{cd}(\mathcal{T}_n) = 2n - 4$ and that $H_{2n-4}(\mathcal{T}_n, \mathbb{Z})$ is infinitely generated) are analogous to Theorems A and C of the present paper. While proofs for corresponding theorems about $\text{GL}(n, \mathbb{Z})$, $\text{Out}(F_n)$, and $\text{Mod}(S)$ often run parallel, at least on a philosophical level, we encounter here a situation where this paradigm does not hold. For instance, the construction of the complex of minimizing cycles in this paper is predicated on the fact that there can be more than one shortest representative for a first homology class in a hyperbolic surface; in a metric graph there is exactly one shortest representative for a first homology class. Also, the proof of the upper bound for $\text{cd}(\mathcal{T}_n)$ does not (in any obvious way) give the correct upper bound for $\text{cd}(\mathcal{I}(S))$.

2. THE COMPLEX OF MINIMIZING CYCLES

After introducing some definitions and terminology, we will define the complex of minimizing cycles for a surface S .

Cycles and curves. A 1-cycle in S is a finite formal sum

$$\sum k_i c_i,$$

where $k_i \in \mathbb{R}$, and each c_i is an oriented simple closed curve in S ; the set $\{c_i : k_i \neq 0\}$ is called the *support*. We say that the 1-cycle is *simple* if the curves of the support are pairwise disjoint, and we say that it is *positive* if each k_i is positive. If we want to emphasize that the k_i are in \mathbb{R} , \mathbb{Q} , or \mathbb{Z} , we can call a cycle *real*, *rational*, or *integral*.

A *basic cycle* is a simple positive 1-cycle $\sum k_i c_i$ with the property that the $[c_i]$ form a linearly independent subset of $H_1(S, \mathbb{R})$. If x is an element of $H_1(S, \mathbb{Z})$, a basic cycle for x (that is, a basic cycle representing x) must be integral.

Let \mathcal{S} be the set of isotopy classes of oriented simple closed curves in S . By taking isotopy classes, there is a natural linear map from the vector space of 1-cycles in S to $\mathbb{R}^{\mathcal{S}}$. Formally, a point of $\mathbb{R}^{\mathcal{S}}$ is a function from \mathcal{S} to \mathbb{R} , and the image of the 1-cycle $\sum k_i c_i$ in $\mathbb{R}^{\mathcal{S}}$ is the function given by

$$s \mapsto \sum_{c_i \in s} k_i.$$

A *multicurve* in the surface S is a nonempty collection of disjoint simple closed curves in S that are homotopically nontrivial and homotopically distinct.

The construction. The first step is to fix an arbitrary nontrivial element x of $H_1(S, \mathbb{Z})$. Next, let \mathcal{M} be the set of isotopy classes of oriented multicurves M in S with the property that each oriented curve of M appears in the support of some basic cycle for x supported in M . The set \mathcal{M} has a natural partial ordering under inclusion (the orientation is important here). As such, we can think of \mathcal{M} as a category where the morphisms are inclusions.

Given $M \in \mathcal{M}$, let P_M be the polytope in $\mathbb{R}^{\mathcal{S}}$ given by the convex hull of the images of the basic cycles for x supported (with orientation) in M . The collection $\{P_M\}$ is a category under inclusion, which we denote \mathcal{F} . The functor given by

$$M \mapsto P_M$$

is an isomorphism from \mathcal{M} to \mathcal{F} . We define the *complex of minimizing cycles* $\mathcal{B}(S)$ as the colimit of this functor:

$$\mathcal{B}(S) = \operatorname{colim}_{M \in \mathcal{M}} \{P_M\}$$

(the choice of x is suppressed in the notation). We can thus regard \mathcal{F} as the category of cells of $\mathcal{B}(S)$.

In different language, $\mathcal{B}(S)$ is the complex obtained from the disjoint union $\coprod P_M$ by identifying faces that are equal in $\mathbb{R}^{\mathcal{S}}$ and endowing the quotient with the weak topology.

Since $\mathcal{I}(S_g)$ acts trivially on $H_1(S_g, \mathbb{Z})$, it acts on $\mathcal{B}(S_g)$. In the next two sections, we show that $\mathcal{B}(S_g)$ is homotopy equivalent to Teichmüller space, and so it is contractible. In particular, we will introduce the notion of a minimizing cycle and explain how $\mathcal{B}(S_g)$ encodes the set of minimizing cycles. An alternate proof of the contractibility of $\mathcal{B}(S)$ is given in Section 5.

Independence of x . Different choices of x in the definition of $\mathcal{B}(S)$ give isomorphic complexes. The mapping class group acts transitively on primitive elements of $H_1(S_g, \mathbb{Z})$, and this action identifies the various complexes associated to primitive vectors. Also, if we replace a vector x with mx , then it follows from Lemma 3.3 below that the map of cycles $\sum k_i c_i \mapsto \sum (mk_i) c_i$ induces an isomorphism of the two complexes of minimizing cycles.

Examples of cells. Consider the picture on the left hand side of Figure 1. Suppose that the homology classes of the curves a , b , and c , are x , y , and $x - y$, respectively. There are exactly two basic cycles for x supported on this multicurve, namely a and $b + c$. Thus, the associated cell of $\mathcal{B}(S_2)$ is an edge.

The three curves in the picture in the middle of Figure 1 are all homologous; say they each represent the homology class x . The basic cycles are simply a , b , and c , and the corresponding cell of $\mathcal{B}(S_4)$ is a 2-simplex.

In the right hand side of the figure, let $[c] + [f] = x$. The basic cycles for x are

$$\{s c + t f + (1 - s)(a + b) + (1 - t)(d + e) : s, t \in [0, 1]\}.$$

Thus, the associated cell is a square. We can think of this example as two independent copies of the first example.




FIGURE 1. Multicurves with their associated cells. The dotted lines indicate the (local) subdivision of $\mathcal{T}(S)$ into chambers (see Section 3.4).

While the examples of Figure 1 are, in a sense, canonical examples, the cells of $\mathcal{B}(S)$ come in a wide variety of shapes. Let a , b , c , d , e , and f be the oriented curves in S_6 shown in the left hand side of Figure 2, and consider, for instance, the homology class $x = [d] + 2[e] + [f]$. There are relations $[a] + [b] + [c] = [e]$ and $[a] + [b] + [d] = [f]$. One can check that the resulting cell is a pentagon, as shown in the right hand side of Figure 2.




FIGURE 2. Curves that give rise to a pentagonal cell of $\mathcal{B}(S_6)$. The picture is drawn in the “ ef -plane”.

Dimension. We now characterize the dimension of a cell of $\mathcal{B}(S)$ and compute the dimension of the complex.

Given $M \in \mathcal{M}$, let $|M|$ denote the number of curves of M , let $D = D(M)$ be the dimension of the span of the images of the curves of M in $H_1(S, \mathbb{R})$, let N be the number of components of $S - M$, and let $B = B(M)$ be the dimension of the associated cell of $\mathcal{B}(S)$.

Lemma 2.1. *For any $M \in \mathcal{M}$, we have $B = |M| - D = N - 1$.*

Proof. There is a natural map from the real vector space spanned by the classes of curves of M to $H_1(S, \mathbb{R})$. The number B is the dimension of the preimage of x ; this is exactly the dimension of the kernel, which is the same as $|M| - D$. The second equality follows from basic algebraic topology. \square

The following proposition is not used anywhere in the paper, but it is a basic fact about the complex of minimizing cycles.

Proposition 2.2. *For $g \geq 2$, the dimension of $\mathcal{B}(S_g)$ is $2g - 3$.*

Proof. By Lemma 2.1, the dimension of the cell of $\mathcal{B}(S_g)$ associated to the oriented multicurve M is one less than the number of connected components of $S_g - M$. Since each component of $S_g - M$ has negative Euler characteristic, the largest possible number of components is $-\chi(S) = 2g - 2$; this is when M is a pants decomposition. Thus, the dimension of a cell of $\mathcal{B}(S_g)$ is at most $2g - 3$. On the other hand, by choosing a pants decomposition of S_g that supports x and consists only of (properly oriented) nonseparating curves, we obtain a cell of dimension $2g - 3$. \square

Simplicial structure. Allen Hatcher has pointed out to us that the complex $\mathcal{B}(S)$ has a natural simplicial structure. So, for instance, in our example of the square cell of $\mathcal{B}(S)$ given above, there is a natural choice of diagonal, subdividing the square into two triangles. The idea is that a point x of $\mathcal{B}(S)$ represented by a cycle c can be translated into a map $S \rightarrow S^1$ in such a way that the components of the complement of a neighborhood of c each map to a point of S^1 (cf. the proof of Lemma 3.3). Varying the distances between these points in S^1 gives the desired simplex containing x .

3. MINIMIZING CYCLES AND THE COMPLEX $\mathcal{B}(S)$

The goal of this section is to give an alternate characterization of $\mathcal{B}(S)$ that is related to Teichmüller space. The idea is to decompose Teichmüller space into regions indexed by the shortest representatives of the homology class x chosen in Section 2.

Teichmüller space. For a surface S with $\chi(S) < 0$ and no boundary, we think of Teichmüller space $\mathcal{T}(S)$ as the space of equivalence classes of pairs $\{(X, \psi)\}$, where X is a complete hyperbolic surface, $\psi : S \rightarrow X$ is a homeomorphism, and $(X, \psi) \sim (X', \psi')$ if $\psi' \circ \psi^{-1}$ is isotopic to an isometry; see [24] for an introduction.

There is a natural action of $\text{Mod}(S)$ on $\mathcal{T}(S)$ given by the formula $[\phi] \cdot (X, \psi) = (X, \psi \circ \phi^{-1})$.

It is a classical theorem that $\mathcal{T}(S)$ is diffeomorphic to $\mathbb{R}^{-3\chi(S)}$ when $\chi(S) < 0$. In particular, and most important for us, $\mathcal{T}(S)$ is contractible.

Lengths of curves. Given a 1-cycle $c = \sum k_i c_i$ in S , and a point $X = (X, \psi)$ in $\mathcal{T}(S)$, we can define the *length of c in X* as

$$\ell_X(c) = \sum |k_i| \ell_X(c_i),$$

where $\ell_X(c_i)$ is the length of the geodesic representative of $\psi(c_i)$ in the hyperbolic metric of X . It is an important and basic fact that the function $\ell_X(c) : \mathcal{T}(S) \rightarrow \mathbb{R}_+$ is continuous.

3.1. Minimizing cycles. We focus our attention on a closed surface S_g . As in the definition of $\mathcal{B}(S_g)$, there is a fixed element x of $H_1(S_g, \mathbb{Z})$, chosen once and for all. Given $X \in \mathcal{T}(S_g)$, we denote by $|x|_{\mathbb{R}}$ the infimum of lengths in X of all real cycles in S_g representing x , and we define a (real) *minimizing cycle* to be a cycle that realizes $|x|_{\mathbb{R}}$.

The goal of this subsection is to prove that minimizing cycles exist (this is not obvious!) and to describe them in terms of multicurves. In the next subsection, we will see that, for various points of $\mathcal{T}(S_g)$, the associated spaces of minimizing cycles are polytopes that exactly correspond to cells of $\mathcal{B}(S_g)$.

In what follows, integral minimizing cycles, rational minimizing cycles, $|x|_{\mathbb{Z}}$, and $|x|_{\mathbb{Q}}$ are defined analogously to the real case.

Integral minimizing cycles. We start by proving the necessary facts about minimizing cycles in the integral case (see also [35]), then proceed to the rational case, and finally the real case.

Lemma 3.1. *Integral minimizing cycles exist; that is, $|x|_{\mathbb{Z}}$ is realized by an integral cycle. There are finitely many integral minimizing cycles.*

Proof. Let c be any integral cycle representing x . There are finitely many geodesics in X with length less than or equal to $\ell_X(c)$, and hence there are finitely many candidates for the support of an integral minimizing cycle. Given any choice of support, there are finitely many choices of integral coefficients that result in a cycle of length less than or equal to $\ell_X(c)$. The lemma follows. \square

Lemma 3.2. *An integral minimizing cycle is simple.*

Proof. Consider an integral cycle that is not simple. We may assume that its self-intersections are transverse and that each curve of the support is a geodesic. If a curve is labeled by an integer k , we replace it with k parallel copies. Then, we resolve all intersections as in Figure 3. The result still represents x . If we choose the parallel copies close enough to the geodesics, then the new cycle obtained after surgery is shorter than the original cycle. \square




FIGURE 3. The resolution of an intersection.

We now want to prove that any two integral minimizing cycles have no transverse intersections. This will follow from the next lemma, which appears, for instance, in work of Thurston [45, Lemma 1].

Lemma 3.3. *Let y be a nontrivial element of $H_1(S_g, \mathbb{Z})$, and let k be a positive integer. Any simple integral 1-cycle in S_g representing ky can be written as the sum of k integral representatives for y .*

Proof. Via Poincaré duality, we can think of ky as an element of $H^1(S_g, \mathbb{Z})$. As such, we obtain a map $f : S_g \rightarrow S^1$ with the properties that $f^*([S^1]) = ky$ and, for any regular value $p \in S^1$, $f^{-1}(p)$ is equal to the given simple chain representing ky . The map f is explicitly constructed as follows. Say $c = \sum k_i c_i$ is a simple 1-cycle representing ky . We thicken each c_i to an annulus $S^1 \times [0, k_i]$ and map each interval $* \times [0, k_i]$ to an interval of length k_i in S^1 (the direction is determined by the orientation of c_i). The complement of the annuli is mapped to a single point. Since ky is an integral class, this gives a well-defined map, uniquely determined up to homotopy.

The singular value of f is the image of the complement of the annuli. The collection of preimages of the regular values of f is a union of disjoint simple closed curves. If we replace each set of parallel curves by a single curve with weight determined by the measure of the image in S^1 , we recover c .

Since the image of $\pi_1(S_g)$ under f_* is $k\mathbb{Z}$, we can lift f as follows:

$$\begin{array}{ccc}
 & & S^1 \\
 & \nearrow \tilde{f} & \downarrow \pi \\
 S_g & \xrightarrow{f} & S^1
 \end{array}$$

where π is multiplication by k . Now $\pi^{-1}(p) = \{p_1, \dots, p_k\}$ is a set of k regular values for \tilde{f} , and by construction, $\bigcup \tilde{f}^{-1}(p_i) = f^{-1}(p)$ is the original representative for ky , and each $\tilde{f}^{-1}(p_i)$ is an integral cycle that represents y . \square

Lemma 3.4. *For $q \in \mathbb{Z}$, we have $|qx|_{\mathbb{Z}} = q|x|_{\mathbb{Z}}$.*

Proof. Let c' be an integral minimizing cycle for qx . By Lemma 3.2, c' is simple, and so by Lemma 3.3, it splits into q representatives of x . Each of these representatives has length at least $|x|_{\mathbb{Z}}$, and so the lemma follows. \square

Lemma 3.5. *Two integral minimizing cycles have no transverse intersections.*

Proof. Suppose that two integral minimizing cycles for x have transverse intersections. By Lemma 3.4, their sum would then be an integral minimizing cycle for $2x$ that is not simple, contradicting Lemma 3.2. \square

Rational minimizing cycles. The next lemma will help bridge the gap between integral minimizing cycles and real minimizing cycles.

Lemma 3.6. *We have $|x|_{\mathbb{Q}} = |x|_{\mathbb{Z}}$; in particular, rational minimizing cycles exist. Also, any rational minimizing cycle is simple, and any two rational minimizing cycles are disjoint.*

Proof. Suppose c is a rational cycle for x with length less than $|x|_{\mathbb{Z}}$. Choose $q \in \mathbb{Z}$ so that qc is integral. Then qc is an integral cycle representing qx with length less than $q|x|_{\mathbb{Z}}$, contradicting Lemma 3.4

Consider any rational minimizing cycle c that is not simple. Choose $q \in \mathbb{Z}$ so that qc is integral. By Lemma 3.4, the cycle qc is then an integral minimizing cycle for qx that is not simple, contradicting Lemma 3.2.

If we have two rational minimizing cycles that intersect, then we can take the “average” of the two cycles to get a rational minimizing cycle for x that is not simple, violating the second conclusion of the lemma. \square

The Borrowing Lemma. The following lemma is the key technical statement that will allow us to relate the complex $\mathcal{B}(S_g)$ to the set of minimizing cycles.

Lemma 3.7 (Borrowing Lemma). *Let $c = \sum_{i=1}^N k_i c_i$ be a real positive cycle representing $x \in H_1(S_g, \mathbb{Z})$. Suppose that (after reindexing) there is a relation*

$$v_1[c_1] + \cdots + v_m[c_m] = v_{m+1}[c_{m+1}] + \cdots + v_{m+n}[c_{m+n}]$$

in $H_1(S_g, \mathbb{R})$, where each v_i is positive. Let

$$L_1 = \sum_{i=1}^m v_i \ell_X(c_i) \quad \text{and} \quad L_2 = \sum_{i=m+1}^{m+n} v_i \ell_X(c_i)$$

and assume (without loss of generality) that $L_1 \leq L_2$. Let δ be a positive real number less than or equal to $\min\{k_i/v_i\}$, where the minimum is taken over $m+1 \leq i \leq m+n$. The cycle

$$c_\delta = \sum_{i=1}^m (k_i + \delta v_i) c_i + \sum_{i=m+1}^{m+n} (k_i - \delta v_i) c_i + \sum_{i=m+n+1}^N k_i c_i$$

is a representative for x . Furthermore, $\ell_X(c_\delta) = \ell_X(c)$ if and only if $L_1 = L_2$ and $\ell_X(c_\delta) < \ell_X(c)$ if and only if $L_1 < L_2$. Finally, the supports of c and c_δ are equal if $\delta < \min\{k_i/v_i\}$.

The proof of the Borrowing Lemma is a simple calculation, which we leave to the reader. As a basic example in genus 2, consider the curves a , b , and c in Figure 4. If the homology classes are x , y , and $x - y$, respectively, and the lengths ℓ_a , ℓ_b , and ℓ_c are close to zero and satisfy $\ell_a = \ell_b + \ell_c$, then the cycle $\epsilon a + (1 - \epsilon)b + (1 - \epsilon)c$ is minimizing for any $\epsilon \in [0, 1]$. This interval's worth of minimizing cycles exactly corresponds to the edge in the left hand side of Figure 1.




FIGURE 4. Curves in S_2 that give rise to “borrowing”.

Real minimizing cycles. Finally, we have the desired result about real minimizing cycles.

Lemma 3.8. *We have $|x|_{\mathbb{R}} = |x|_{\mathbb{Q}} = |x|_{\mathbb{Z}}$; in particular, real minimizing cycles exist. Also, any real minimizing cycle is simple, and any two real minimizing cycles are disjoint.*

Proof. Let $c = \sum k_i c_i$ be a nonrational real cycle. Since x is rational, it follows that the set of c_i with k_i irrational are linearly dependent. An application of the Borrowing Lemma allows us to reduce the number of irrational coefficients without

increasing the length of the representative. In the end, we get a rational cycle. Thus, $|x|_{\mathbb{R}} = |x|_{\mathbb{Q}}$, and these are equal to $|x|_{\mathbb{Z}}$ by Lemma 3.6.

If c is a real cycle that is not simple, we can, as above, use the Borrowing Lemma to get a rational cycle of the same (or lesser) length that is also not simple. Since rational minimizing cycles are simple (Lemma 3.6), it follows that c is not a real minimizing cycle.

The proof of the last statement follows similarly. □

3.2. The space of minimizing cycles. The goal of this subsection is to describe the parameter space of minimizing cycles for any one point of $\mathcal{T}(S_g)$. We will see that each such parameter space corresponds to a cell of $\mathcal{B}(S_g)$. As above, $x \in H_1(S_g, \mathbb{Z})$ and $X \in \mathcal{T}(S_g)$ are fixed.

Minimizing multicurves. A consequence of Lemma 3.8 is that there is a canonical multicurve $M = M_X$ in S_g , called the *minimizing multicurve*, which is the smallest multicurve with the property that every minimizing cycle for x is supported on M . The minimizing multicurve M is obtained by taking the union of the supports of all minimizing cycles for x . Note that M is a union of geodesics in X .

Lemma 3.9. *There is a unique orientation of the minimizing multicurve $M = M_X$ so that each minimizing cycle for x is a nonnegative linear combination of curves of M .*

Proof. Suppose M is the collection of oriented curves $\{c_1, \dots, c_n\}$ and that $c = \sum k_i c_i$ and $c' = \sum k'_i c_i$ are minimizing cycles for x with $k'_i < 0 < k_i$ for some particular i . In this case, $c + c'$ is a representative for $2x$ with length strictly less than twice that of c . This contradicts Lemma 3.4. □

As a consequence of the lemma, it makes sense to focus our attention on positive minimizing cycles for x .

Borrowing relations. Let $M = \{c_i\}$ be an oriented multicurve. With the Borrowing Lemma in mind, we call any relation amongst the $[c_i]$ in $H_1(S_g, \mathbb{Z})$ a *borrowing relation*. If we cut S_g along M , then each resulting connected component gives rise to a borrowing relation, which we call a *subsurface borrowing relation*.

Lemma 3.10. *Let $X \in \mathcal{T}(S_g)$, and let $M = \{c_i\}$ be an oriented multicurve that is the union of (oriented) supports of positive cycles that represent x . The following are equivalent.*

- (1) *All positive cycles for x supported in M have the same length.*
- (2) *For every subsurface borrowing relation $\sum v_i [c_i] = 0$, the $\ell_X(c_i)$ satisfy $\sum v_i \ell_X(c_i) = 0$.*

In particular, if M is a minimizing multicurve, then both conditions are satisfied.

Proof. By the Borrowing Lemma, statement (1) is equivalent to the statement: (3) for every borrowing relation $\sum v_i [c_i] = 0$, the $\ell_X(c_i)$ satisfy $\sum v_i \ell_X(c_i) = 0$. By basic algebraic topology, any borrowing relation between the $[c_i]$ is a consequence of the subsurface borrowing relations between the $[c_i]$. Hence (3) is equivalent to statement (2). □

Borrowing cells. We now explain how the minimizing multicurve $M = M_X$ gives a natural way to parameterize the space of minimizing cycles in X . As in Lemma 3.9, we orient M so that all minimizing cycles are positive.

As in Section 2, $D = D(M)$ is the dimension of the span of M in $H_1(S_g, \mathbb{R})$ and $B = B(M)$ is $|M| - D$. There is a natural map $\mathbb{R}^{|M|} \rightarrow H_1(S_g, \mathbb{R})$, and the preimage of the subspace spanned by x is a B -dimensional affine subspace V_X of $\mathbb{R}^{|M|}$. By taking lengths of cycles in X , we get a map $L_X : V_X \rightarrow \mathbb{R}_+$.

We can think of the minset P_X of L_X in V_X as the parameter space of minimizing cycles for x in X . As P_X is the minset of a sum of convex functions, it is a convex set P_X , which we call a *borrowing cell*. Since each of these (finitely many) functions is proper, it follows that P_X is a compact set. Finally, since each of the functions is the absolute value of a linear function, it follows that P_X is a polytope (it is cut out by the zero sets of the individual functions).

Lemma 3.11. *Each borrowing cell P_X is a compact convex polytope in V_X .*

We now prove that the borrowing cell P_X is the convex hull of its basic cycles for x . This gives the key connection between minimizing cycles and the complex $\mathcal{B}(S_g)$. Recall from Section 2 that P_M is the cell of $\mathcal{B}(S_g)$ associated to a multicurve M .

Lemma 3.12. *There is a bijective correspondence between the set of basic cycles for x that are supported in $M = M_X$ and the set of vertices of P_X . In particular, $P_X = P_{M_X}$.*

Proof. We use the following characterization: a point of a polytope in Euclidean space is a vertex if and only if every open line segment through that point leaves the polytope.

First, let p be a point of P_X corresponding to a basic minimizing cycle. If we move away from p in any direction of V_X , some coordinate must change from zero to nonzero, and it follows that, in any open line segment of V_X through p , there must be a point with at least one negative coordinate. By Lemma 3.9, such points are not contained in P_X , and so p is a vertex.

Next, let p be a point of P_X that corresponds to a minimizing cycle that is not basic. By the Borrowing Lemma, there is an open line segment of V_X through p so that all points on the line segment have positive coordinates. Since M is a minimizing cycle, it follows from the last sentence of Lemma 3.10 that statement (1) of the same lemma holds, and we see that the line segment is contained in P_X . Thus, p is not a vertex. \square

The reader is invited to revisit the examples of Section 2, and reinterpret them as borrowing cells in the context of minimizing cycles.

Recall from Section 2 that \mathcal{M} is the category of oriented multicurves that are made up of basic cycles for x , and \mathcal{F} is the category of cells of $\mathcal{B}(S_g)$. The set of borrowing cells also forms a category under inclusion if we view each borrowing cell as a subset of \mathbb{R}^S .

Lemma 3.13. *The following categories are isomorphic: \mathcal{M} , \mathcal{F} , and the category of borrowing cells.*

Proof. The equivalence of \mathcal{M} and \mathcal{F} was already stated in Section 2. Also, by Lemma 3.12, the functor $P_X \mapsto M_X$ is an isomorphism from the category of borrowing cells to its image in the category \mathcal{M} . It remains to show that this functor is surjective.

In other words, we need to produce, for a given $M \in \mathcal{M}$, a marked hyperbolic surface $X \in \mathcal{T}(S_g)$ with $M_X = M$. This is done by choosing the lengths of the curves of M to be much smaller than all other curves in X , and inductively choosing the lengths of the curves of M to satisfy the subsurface borrowing relations, as in Lemma 3.10. \square

3.3. Chambers. To each borrowing cell (equivalently, to each element of \mathcal{M}), we now associate a subset of $\mathcal{T}(S_g)$, which we call a “chamber”. We will show that these chambers are contractible (Proposition 3.14), and that the cover of $\mathcal{T}(S_g)$ by chambers is, in a sense, dual to the cover of $\mathcal{B}(S_g)$ by cells. Since $\mathcal{T}(S_g)$ is contractible, it will follow that $\mathcal{B}(S_g)$ is contractible.

Let $M \in \mathcal{M}$. We define the *closed chamber* associated to M as

$$Y_M = \{X \in \mathcal{T}(S_g) : M \subseteq M_X\}.$$

The following proposition is proved in Section 4 and is the key technical part of the paper.

Proposition 3.14. *Each closed chamber Y_M is contractible.*

The next lemma is used in the proof of Proposition 3.14. See the appendix of [8] for an introduction to manifolds with corners. For the statement, we define the *open chamber* associated to the oriented multicurve M as follows:

$$Y_M^o = \{X \in \mathcal{T}(S_g) : M = M_X\}.$$

The proof of the lemma uses the Fenchel–Nielsen coordinates for $\mathcal{T}(S_g)$ (see [24, Section 3.2]) and the following basic fact from Teichmüller theory.

Fact 3.15. Let $X_0 \in \mathcal{T}(S_g)$, and $D > 0$. There is a neighborhood U of X_0 so that the set $\{c \in \mathcal{S} : \ell_X(c) < D \text{ for some } X \in U\}$ is finite.

Lemma 3.16. *Each closed chamber Y_M is a manifold with corners. Its interior is the open chamber Y_M^o .*

Proof. Let $X \in Y_M$. Choose a pants decomposition for S_g that contains the minimizing multicurve for X (this multicurve necessarily contains M), and consider the resulting Fenchel–Nielsen coordinates on $\mathcal{T}(S_g)$. There is a finite set of linear inequalities among the length coordinates so that, in a neighborhood U of X , a point is in Y_M if and only if those inequalities are satisfied. Since the solution set of a system of linear inequalities in \mathbb{R}^n is a manifold with corners, and since Fenchel–Nielsen coordinates form a smooth chart for $\mathcal{T}(S_g)$, it follows that Y_M is a manifold with corners.

In the previous paragraph, the locus of points in \mathbb{R}^n where all of the inequalities are strict is exactly where M is the minimizing multicurve; but this is the definition of Y_M^o . \square

We now relate the partial orderings on multicurves and closed chambers.

Lemma 3.17. *$Y_M \supseteq Y_{M'}$ if and only if $M \subseteq M'$.*

Proof. By the definition of closed chambers, the statement $Y_M \supseteq Y_{M'}$ is equivalent to the statement that, for any $X \in \mathcal{T}(S_g)$, if $M_X \supseteq M'$, then $M_X \supseteq M$. The latter is equivalent to the inclusion $M \subseteq M'$. \square

Given two multicurves M and M' with no transverse intersections, the multicurve $M \cup M'$ is defined by taking the union of the two multicurves and replacing each pair of isotopic curves with a single curve.

Corollary 3.18. *The intersection of two closed chambers Y_M and $Y_{M'}$ is nonempty if and only if the multicurve $M \cup M'$ is defined and has an orientation compatible with those of M and M' . In this case, $Y_M \cap Y_{M'} = Y_{M \cup M'}$.*

From the above corollary we can deduce that the category of closed chambers (under inclusion) and the opposite of the category \mathcal{M} are isomorphic. We will not use this fact directly, but the contravariance belies the main difficulty in the proof of Theorem E in Section 3.4.

Corollary 3.19. *The nerve of the cover of $\mathcal{T}(S_g)$ by the closed chambers $\{Y_M\}$ is finite dimensional.*

Proof. Suppose $\{Y_{M_i}\}$ is a collection of closed chambers with nontrivial intersection. By Corollary 3.18, this intersection is a closed chamber Y_M , and each M_i is a submulticurve of M . Now, M consists of at most $3g-3$ curves, and so the set $\{Y_{M_i}\}$ contains at most 2^{3g-3} elements. This gives an upper bound for the dimension of the nerve. \square

3.4. The complex of minimizing cycles is contractible. In this section, we prove Theorem E by making a comparison between the cover of $\mathcal{B}(S_g)$ by cells and the cover of $\mathcal{T}(S_g)$ by chambers. Roughly, since both chambers and cells of $\mathcal{B}(S_g)$ are contractible, and since the gluing patterns are the same (up to contravariance), it will follow that $\mathcal{B}(S_g)$ is homotopy equivalent to $\mathcal{T}(S_g)$, and so $\mathcal{B}(S_g)$ is contractible.

Proof of Theorem E. Let $\mathcal{B}'(S_g)$ denote the barycentric subdivision of $\mathcal{B}(S_g)$; it suffices to show that $\mathcal{B}'(S_g)$ is contractible. For a vertex of $\mathcal{B}(S_g)$ corresponding to a multicurve M , denote by C_M the star of the corresponding vertex in $\mathcal{B}'(S_g)$.

The proof is in three steps.

- (1) The geometric realization of the nerve of the cover of $\mathcal{T}(S_g)$ by closed chambers is homotopy equivalent to $\mathcal{T}(S_g)$.
- (2) The geometric realization of the nerve of the cover of $\mathcal{B}'(S_g)$ by the $\{C_M\}$ is homotopy equivalent to $\mathcal{B}'(S_g)$.
- (3) The two nerves are the same.

Step 1. It is a special case of a result of Borel–Serre that if a topological space is covered by a set of closed, contractible manifolds with corners, and if the cover is closed under taking nonempty intersections, and the dimension of the nerve of the cover is finite, then the nerve is homotopy equivalent to the original space [8, Theorems 8.2.1, 8.3.1]. Thus, it suffices to apply Proposition 3.14, Lemma 3.16, Corollary 3.18, and Corollary 3.19.

Step 2. We apply the same theorem of Borel–Serre. In general, the closed star of a vertex of a simplicial complex is contractible ($\mathcal{B}(S_g)$ is not simplicial, but its barycentric subdivision is). In the category of simplicial complexes, the phrase “closed manifold with corners” used in Step 1 can be replaced with “subcomplex”

(both types of spaces are examples of absolute neighborhood retracts). That the dimension of the nerve is finite follows from Step 3 (which is independent of the other steps). Thus, their theorem applies.

Step 3. A collection $\{C_{M_i}\}$ has nonempty intersection if and only if the vertices M_i are vertices of a common cell of $\mathcal{B}(S_g)$. By Lemma 3.13, this is equivalent to the condition that the M_i are vertices of some common borrowing cell, that is, if there is a hyperbolic structure X with $M_i \subseteq M_X$ for all i . But this is the same as the condition that the intersection of the Y_{M_i} is nonempty (Corollary 3.18).

Combining Steps 1, 2, and 3, plus the fact that $\mathcal{T}(S_g)$ is contractible, we see that $\mathcal{B}'(S_g)$, and hence $\mathcal{B}(S_g)$, is contractible. \square

4. CHAMBERS ARE CONTRACTIBLE

In this section we prove Proposition 3.14, which says that each closed chamber Y_M is contractible. We take x to be the element of $H_1(S_g, \mathbb{Z})$ chosen in Section 2. As in Section 3, Y_M is the closed chamber associated to some multicurve M whose isotopy class is in \mathcal{M} .

4.1. The thick part of Teichmüller space is contractible. The goal is to reduce Proposition 3.14 to the following theorem of Ivanov [25, Theorem 3]. The ϵ -*thick part* of Teichmüller space $\mathcal{T}_\epsilon(S)$ is the subset consisting of hyperbolic surfaces where every nontrivial closed geodesic that is not a boundary component has length at least ϵ (in the case of a surface with boundary, we define Teichmüller space as the space of marked hyperbolic surfaces where each boundary is a geodesic of fixed length).

Also, we make use of the Margulis constant $\epsilon_{\mathbb{H}^2}$ for the hyperbolic plane. This number has the property that, in any hyperbolic surface, two essential nonisotopic curves of length less than $\epsilon_{\mathbb{H}^2}$ are disjoint.

Theorem 4.1 (Ivanov). *Let S be any compact surface and let $\epsilon \in (0, \epsilon_{\mathbb{H}^2})$. The space $\mathcal{T}_\epsilon(S)$ is contractible.*

For any $\epsilon \in (0, \epsilon_{\mathbb{H}^2})$ and $p \in \mathbb{R}_+^{|M|}$, we define the subset $\mathcal{A}_{\epsilon,p}$ of $\mathcal{T}(S_g)$ to be the set of points with the following properties.

- (1) Each curve in S_g that is disjoint from the curves of M and is not isotopic to a curve of M has length at least ϵ .
- (2) The lengths of the curves of M are given by p .

Lemma 4.2. *For any $\epsilon < \epsilon_{\mathbb{H}^2}$, there is a choice of p so that $\mathcal{A}_{\epsilon,p}$ is a subset of the open chamber Y_M^o .*

Proof. Let $\epsilon < \epsilon_{\mathbb{H}^2}$ be given. We choose $p \in \mathbb{R}_+^{|M|}$ so that every positive cycle for x supported (with orientation) in M has the same length less than ϵ (take any point in Y_M^o , and shrink the curves of M by the same factor until this condition holds).

In $\mathcal{A}_{\epsilon,p}$, a basic minimizing cycle for x cannot have support on any curve that is disjoint from (and not isotopic to) M , since all such curves have length at least ϵ (by the definition of $\mathcal{A}_{\epsilon,p}$). Also, no geodesic intersecting the geodesic representative of M transversely can be in the support of a basic minimizing cycle, since all such curves have length greater than ϵ by the hypothesis on ϵ . Thus, $\mathcal{A}_{\epsilon,p} \subset Y_M^o$. \square

Lemma 4.3. *For any $\epsilon < \epsilon_{\mathbb{H}^2}$ and any p , the space $\mathcal{A}_{\epsilon,p}$ is contractible.*

Proof. Take a pants decomposition P of S_g containing M , and consider the associated Fenchel–Nielsen coordinates.

Let $\{R_i\}$ be the set of connected components of $S_g - M$. There is a natural map

$$\mathcal{A}_{\epsilon,p} \rightarrow \mathcal{T}_\epsilon(R_i)$$

for each i , since each R_i has an induced marking and an induced hyperbolic structure (we make sure to cut along the geodesic representative of M).

We also consider the $|M|$ maps

$$\mathcal{A}_{\epsilon,p} \rightarrow \mathbb{R}$$

obtained by taking the twist parameters at the curves of M .

Combining these maps, we get a bijective map

$$\mathcal{A}_{\epsilon,p} \rightarrow \prod \mathcal{T}_\epsilon(R_i) \times \mathbb{R}^{|M|}.$$

This map is a homeomorphism; indeed, there is a continuous inverse, obtained by gluing the pieces back together according to the prescribed twist parameters. Thus, applying Theorem 4.1, we see that $\mathcal{A}_{\epsilon,p}$ is contractible. \square

4.2. Chamber flow. The goal of this section is to prove the following technical statement, which, combined with Lemmas 4.2 and 4.3, completes the proof that closed chambers are contractible (Proposition 3.14).

Lemma 4.4. *Given any compact subset K of the closed chamber Y_M , there is a choice of $\mathcal{A}_{\epsilon,p} \subset Y_M^\circ$, and a deformation of a subset of Y_M that takes K into $\mathcal{A}_{\epsilon,p}$.*

For the proof, we need two lemmas from Teichmüller theory. For a pants decomposition P of S_g , let FN_P be the map $\mathbb{R}^{6g-6} \rightarrow \mathcal{T}(S_g)$ corresponding to the Fenchel–Nielsen coordinates adapted to P , and let $(FN_P)_\star$ be the differential. In the statement of the lemma, we endow \mathbb{R}^{6g-6} with the Euclidean metric, and $\mathcal{T}(S_g)$ with the Teichmüller metric.

Lemma 4.5. *Let $D_{\min}, D_{\max} \in \mathbb{R}_+$. There exists a constant $C = C(g, D_{\min}, D_{\max})$ so that if P is any pants decomposition of S_g , and*

$$W = \{X \in \mathcal{T}(S_g) : D_{\min} \leq \ell_X(c) \leq D_{\max} \text{ for all } c \in P\},$$

then $(FN_P)_\star$ is C -Lipschitz on $FN_P^{-1}(W)$.

Proof. Let P be given. The subset $FN_P^{-1}(W)$ of \mathbb{R}^{6g-6} is the infinite “strip”

$$[D_{\min}, D_{\max}]^{3g-3} \times \mathbb{R}^{3g-3},$$

that is, all length parameters are in $[D_{\min}, D_{\max}]$, and twist parameters are unconstrained. The free abelian group $G(P)$ generated by Dehn twists in the curves of P acts on W , since $G(P)$ preserves the lengths of the curves of P . Moreover, this action is by isometries (in general, $\text{Mod}(S)$ acts on $\mathcal{T}(S)$ by isometries).

In the Fenchel–Nielsen coordinates, the Dehn twist about a curve of P is translation by 2π in the corresponding twist coordinate. Therefore, the quotient of $FN_P^{-1}(W)$ by $G(P)$ is

$$[D_{\min}, D_{\max}]^{3g-3} \times (S^1)^{3g-3}.$$

Again, this action is by isometries, so the stretch factor of $(FN_P)_\star$ is determined on this quotient. As this quotient is compact, and FN_P is smooth, it follows that $(FN_P)_\star$ is Lipschitz on $FN_P^{-1}(W)$.

To get the Lipschitz constant C which only depends on g and the constants D_{\min} and D_{\max} (and not on the pants decomposition P), we note that there are only finitely many topological types of pants decompositions on S_g . Since $\text{Mod}(S_g)$ acts by isometries on $\mathcal{T}(S_g)$, and since $\text{Mod}(S_g)$ preserves the infinite strips $[D_{\min}, D_{\max}]^{3g-3} \times \mathbb{R}^{3g-3}$, the Lipschitz constant is the same for two pants decompositions that are topologically equivalent. Therefore, the constant $C = C(g, D_{\min}, D_{\max})$ is the maximum of finitely many Lipschitz constants obtained as above. \square

The next lemma is the technical statement that implies the existence of the Bers constant, that is, the universal constant L_g with the property that every hyperbolic surface of genus g has a pants decomposition of total length at most L_g (cf. the inductive step of [10, Theorem 5.2.3]).

Lemma 4.6. *Let $g \geq 2$ and $L > 0$. There is a constant $D = D(g, L)$ so that, if M is a geodesic multicurve of $X \in \mathcal{T}(S_g)$, where each curve of M has length at most L , then there is a pants decomposition P containing M so that each curve of P has length less than $D(g, L)$.*

We now construct the deformation of K into $\mathcal{A}_{\epsilon, p}$. The argument is modelled on Ivanov's proof of Theorem 4.1.

Proof of Lemma 4.4. We will define a smooth vector field \mathcal{V} on an open subset of Y_M , and the resulting flow will give the required deformation. The basic idea is to use Fenchel–Nielsen coordinates to define a vector field that, at a given point, points in a direction where the following conditions hold.

- The curves of M are getting shorter, and their lengths are approaching the point $p \in \mathbb{R}_+^{|M|}$.
- Any curve that is not in M , is disjoint from M , and has length ϵ stays the same length; in particular, no curve outside of M gets shorter than ϵ (we will choose $\epsilon > 0$ so that all curves begin with length greater than ϵ).
- Any minimizing cycle for x not supported in M has at least one curve that stays the same length (in particular, by the first condition, minimizing cycles not supported in M immediately become nonminimizing).

In short, vectors should point towards the interior of Y_M (third condition), and in particular towards $\mathcal{A}_{\epsilon, p}$ (first two conditions). To accomplish this, there are three steps:

- Step 1. Make a choice of $\mathcal{A}_{\epsilon, p}$.
- Step 2. Define the vector field \mathcal{V} locally.
- Step 3. Check that \mathcal{V} has the desired properties.

Step 1. Picking an $\mathcal{A}_{\epsilon, p}$ involves choosing ϵ and p . Let L_K be the length of the longest minimizing cycle for x in K . We choose $\epsilon < \epsilon_{\mathbb{H}^2}$ small enough so that the following two conditions are satisfied.

- Every nontrivial geodesic in each surface of K has length at least ϵ .
- If c is a geodesic that crosses another geodesic of length at most ϵ , then c has length greater than L_K .

To see that both conditions can be satisfied for ϵ small enough, we use the fact that the length of a particular curve in S is a continuous function on $\mathcal{T}(S)$, and for the second condition, we apply the collar lemma, which states that, in any

hyperbolic surface, as the length of a geodesic d tends to zero, the length of the shortest geodesic intersecting d transversely tends to infinity.

The point of the first condition on ϵ is that, in order to prevent curves from getting shorter than ϵ , they must be longer than ϵ to begin with. The second condition implies that, in K , and at any point of Y_M where the length of a minimizing cycle is at most L_K , curves of length ϵ (or shorter) have trivial geometric intersection with the curves of M .

In the proof of Lemma 4.2, the choice of p only depended on the choice of ϵ and on M . We choose p in the same way here, using the ϵ defined here. As a result, we have $\mathcal{A}_{\epsilon,p} \subset Y_M^o$.

Step 2. We will define the vector field \mathcal{V} on open sets indexed by pants decompositions. First, let D_{\max} be the constant $D(g, L_K)$ from Lemma 4.6, where L_K is the constant from Step 1. Note that $D_{\max} > L_K$. Choose D_{\min} to be smaller than the smallest coordinate of p in $\mathbb{R}^{|M|}$. Note that $D_{\min} < |p| < \epsilon$, and, in particular, D_{\min} is smaller than the length of any curve in any surface of K .

Let \mathcal{P} be the set of pants decompositions of S_g that contain M . For $P \in \mathcal{P}$, let U_P be the open subset of Y_M given by

$$U_P = \{X \in Y_M \quad : \quad \begin{array}{l} \text{if } X \in Y_{M'}, \text{ then } M' \subseteq P, \\ \text{if } \ell_X(c) \leq \epsilon, \text{ then } c \in P, \\ D_{\min} < \ell_X(c) < D_{\max} \text{ for all } c \in P \}. \end{array}$$

The second condition is indeed an open condition: restate it as $c \notin P \Rightarrow \ell_X(c) > \epsilon$ and apply Fact 3.15.

Let $U = \bigcup_{P \in \mathcal{P}} U_P$. The compact set K is contained in U ; this follows from the choices of D_{\max} and D_{\min} , Lemma 4.6, and the fact that there are no curves of length less than or equal to ϵ in surfaces of K .

For a given pants decomposition P , we now define a vector field \mathcal{V}_P on U_P . Let $d : \mathbb{R}^{|M|} \rightarrow \mathbb{R}_{\geq 0}$ be the square of the Euclidean distance from the point $p \in \mathbb{R}^{|M|}$. Write \mathbb{R}^{6g-6} as $\mathbb{R}^{|M|} \times \mathbb{R}^{6g-6-|M|}$, and let \mathcal{V}'_P be the vector field on $FN_{\overline{P}}^{-1}(U_P) \subset \mathbb{R}^{6g-6}$ given by $(-\nabla d, 0)$. We then define \mathcal{V}_P on U_P via $\mathcal{V}_P = (FN_P)_*(\mathcal{V}'_P)$.

Consider the map

$$\Phi_P : U_P \rightarrow \mathbb{R}^{|M|}$$

obtained by taking the lengths of the curves of M . By construction, $\Phi_P^*(\mathcal{V}_P) = -\nabla d$.

There is a partition of unity subordinate to the cover of U by the U_P . We use this to combine the local vector fields \mathcal{V}_P into a global vector field \mathcal{V} defined on all of U . Since \mathcal{V} is obtained by averaging the \mathcal{V}_P , we still have $\Phi^*(\mathcal{V}) = -\nabla d$, where $\Phi : U \rightarrow \mathbb{R}^{|M|}$ is the map that records the lengths of the curves of M . In other words, even after using the partition of unity, we still know what happens to the lengths of the curves of M along flow lines of \mathcal{V} .

Step 3. We will prove the following statements about \mathcal{V} .

- (1) All flow lines starting in K stay in Y_M .
- (2) All flow lines starting in K stay in U .
- (3) There is a universal constant C so that $\|v\| \leq C\|\Phi^*(v)\|$ for all v .
- (4) All flow lines starting in K converge as time goes to infinity.

The first three statements imply that the flow induced by \mathcal{V} is defined for all positive time. From the fourth statement (plus the definition of \mathcal{V}), we see that, at

infinity, all points of K end up in $\mathcal{A}_{\epsilon,p}$. The resulting map $K \rightarrow \mathcal{A}_{\epsilon,p}$ is automatically continuous. As $\mathcal{A}_{\epsilon,p}$ is contractible (Lemma 4.3), the lemma will follow.

To prove statement (1), there are two things to show: the positive cycles for x supported in M always have equal length, and no other cycles for x get shorter than those. The first holds because the subset of $\mathbb{R}^{|M|}$ where the positive cycles for x have the same length is convex (by Lemma 3.10 it is cut out by the linear equations coming from the subsurface borrowing relations); since we are flowing along straight lines towards p , we stay in this convex set. For the second, we note that, by construction, all vectors on the boundary of Y_M point into the interior of Y_M , i.e. the open chamber Y_M^o .

For statement (2), let X be a point on a flow line that emanates from K . By the construction of \mathcal{V} , the image of Y_M under the flow induced by \mathcal{V} is contained in Y_M^o for all positive time. Therefore, since we already showed that $K \subset U$, we may assume that $X \in Y_M^o$. We need to argue that there is a pants decomposition P of X so that $X \in U_P$. To construct P , we start with M , add all curves of length at most ϵ , and (if needed) complete to a pants decomposition using curves of length at most D_{\max} . We now check that $X \in U_P$; the only nontrivial point is the condition $D_{\min} < \ell_X(c) < D_{\max}$ for all $c \in P$. Each $\ell_X(c)$ is less than D_{\max} by the definition of D_{\max} and the fact that the lengths of the curves of M get shorter along flow lines. Each $\ell_X(c)$ is greater than D_{\min} because each curve of M stays longer than $|p| > D_{\min}$ along flow lines, and curves disjoint from M do not get shorter than $\epsilon > D_{\min}$ along flow lines.

For statement (3), we apply Lemma 4.5. By the definition of U_P , for any $P \in \mathcal{P}$, and any curve c of P , we have $D_{\min} < \ell_X(c) < D_{\max}$. Let $C = C(g, D_{\min}, D_{\max})$ be the constant given by Lemma 4.5. By the lemma, each map $FN_P : FN_P^{-1}(U_P) \rightarrow U_P$ is C -Lipschitz. In other words, the length of a vector of \mathcal{V}_P is at most C times the length of the corresponding vector of \mathcal{V}'_P . The latter is the same as the length of the corresponding vector of $-\nabla d$. Since \mathcal{V} is obtained from the \mathcal{V}_P by averaging, it follows that each vector of \mathcal{V} has length at most C times the length of the corresponding vector of $-\nabla d$ (the correspondence is given by Φ^*), which is what we wanted to show.

From statement (3), it follows that if γ is an integral path in \mathbb{R}^{6g-6} of length L , then the lift of this path to $\mathcal{T}(S_g)$ has length at most CL . Therefore, since flow lines of \mathcal{V}'_P in \mathbb{R}^{6g-6} all converge, the corresponding flow lines in $\mathcal{T}(S_g)$ also converge (statement (4)). This completes the proof. \square

5. A SURGICAL PROOF OF THE CONTRACTIBILITY OF $\mathcal{B}(S)$

In this section we give a short proof of the contractibility of the complex of minimizing cycles (Theorem E). Our proof is inspired by Hatcher’s proof of the contractibility of the arc complex for a surface [23]. One advantage over the proof given in Sections 3 and 4 is that the proof here more easily covers the case where S is not closed.

As in Section 2, let S be a surface, let x be a nontrivial element of $H_1(S, \mathbb{Z})$, and let $\mathcal{B}(S)$ be the associated complex of minimizing cycles.

An auxiliary complex. Let $\widetilde{\mathcal{M}}$ be the set of isotopy classes of oriented multi-curves M in S with the property that each oriented curve of M lies in the support of some 1-cycle that represents x and is supported in M ; the key is that homologically

trivial cycles, such as separating curves, are allowed. If in the definition of $\mathcal{B}(S)$ we replace \mathcal{M} with $\widetilde{\mathcal{M}}$, we get a space, which we denote by $\widetilde{\mathcal{B}}(S)$, that contains $\mathcal{B}(S)$. The space $\widetilde{\mathcal{B}}(S)$ has the structure of a cell complex, except that some of the “cells” are not compact.

A characterization of $\mathcal{B}(S)$. Below, we give a new characterization of points of $\mathcal{B}(S)$. For a positive 1-cycle $c = \sum k_i c_i$, an equation of the form $[c_{i_1}] + \cdots + [c_{i_n}] = 0$ ($n > 0$) is called a *one-sided relation*, and we may say that c *satisfies* a one-sided relation.

Lemma 5.1. *Let c be a 1-cycle in S . The following are equivalent.*

- (1) c represents a point of $\mathcal{B}(S)$.
- (2) c is a minimizing cycle for x .
- (3) c satisfies no one-sided relations.

Proof. The equivalence of (1) and (2) is the content of Lemma 3.12 above.

That (2) implies (3) is straightforward: if c satisfies a one-sided relation, say, $[c_1] + \cdots + [c_n] = 0$, then, by the Borrowing Lemma, we may shorten c by subtracting ϵ from the coefficients of c_1, \dots, c_n , where ϵ is any positive number smaller than the minimum of $\{k_1, \dots, k_n\}$.

To show that (3) implies (1), we suppose that c satisfies no one-sided relations. The goal is to show that there is a basic cycle for x supported in $\{c_i\}$ with some fixed c_{i_0} in the support. If c is not already basic, that means $S - \bigcup c_i$ is disconnected. If S' is one component, then the c_i that form the boundary of S' satisfy a relation of the form

$$[c_1] + \cdots + [c_p] = [c_{p+1}] + \cdots + [c_{p+q}]$$

with $p, q > 0$. Suppose without loss of generality that $i_0 > p$ and that k_1 is the minimum of $\{k_1, \dots, k_p\}$. As in the Borrowing Lemma, the following is a new simple positive 1-cycle that represents x :

$$\begin{aligned} c' = & ((k_2 - k_1)c_2 + \cdots + (k_p - k_1)c_p) \\ & + ((k_{p+1} + k_1)c_{p+1} + \cdots + (k_{p+q} + k_1)c_{p+q}) + \sum_{i>p+q} k_i c_i. \end{aligned}$$

The 1-cycle c' is a simple positive 1-cycle that contains c_{i_0} in its support. What is more, $S - c'$ has one fewer connected component than $S - c$. By induction, we can find a 1-cycle representing x that is supported in $\{c_i\}$ and that is nonseparating in S . Such a 1-cycle is basic. \square

Surgery on 1-cycles. Suppose that d and d' are two 1-cycles in S representing points of $\mathcal{B}(S)$, assume that d and d' have the minimal number of intersections in their isotopy classes, and let c be the 1-cycle $td + (1 - t)d'$, where $t \in [0, 1]$; note that c represents the class x . In general, c is not simple. We now explain how to do surgery to convert c into a simple 1-cycle $\text{Surger}(c)$, which is canonical up to isotopy.

By Poincaré Duality, the homology class x corresponds to a cohomology class, and hence to a homotopy class of based maps from S to the circle S^1 . We will use the cycle c to construct an explicit representative φ of this homotopy class.

As above, let $c = \sum k_i c_i$. We thicken each c_i to a band (annulus) $A_i = S^1 \times [0, k_i]$ in such a way that c_i is a coordinate circle. We choose the A_i and coordinates (θ_i, t_i) so that the following conditions hold (cf. the left hand side of Figure 5).

- (1) Each c_i is oriented in the positive θ direction.
- (2) Any two curves $S^1 \times *$ are transverse.

We now define a real-valued function μ_i on the set of arcs in A_i . For an arc $\alpha : [0, 1] \rightarrow A_i$, we define $\mu_i(\alpha) = \int_{\alpha} dt_i$. We get a function μ_i on arcs in S by adding the μ_i -values of the components contained in A_i . Finally, we obtain a real-valued function μ on arcs in S given by

$$\mu(\alpha) = \sum \mu_i(\alpha).$$

For a closed loop α , we see that $\mu(\alpha)$ is the same as the algebraic intersection of α with c .

We are now ready to define the map $\varphi : S \rightarrow S^1$. Let p be an arbitrarily chosen basepoint of S , let q be an arbitrary point of S , and let $\alpha_q : [0, 1] \rightarrow S$ be an arbitrary path from p to q . Thinking of S^1 as \mathbb{R}/\mathbb{Z} we define $\varphi(q)$ as the fractional part of $\mu(\alpha_q)$. Since x is an integral homology class, this map is well defined.

We now use the map $\varphi : S \rightarrow S^1$ to define the simple 1-cycle $\text{Surger}(c)$. The preimage of each regular value of φ in S^1 is a 1-manifold, that is, a collection of pairwise disjoint curves in S . If we consider the union of the preimages of regular values in S^1 all at once, we see some number of parallel families of curves. Each family has a well-defined thickness, coming from the measure on the circle (which in turn comes from the standard measure on \mathbb{R}). We replace each family of curves with a single weighted curve, where the weight is the total thickness. Also, we discard any curves that are homotopic to a point. The resulting 1-chain is the desired simple 1-chain $\text{Surger}(c)$. From the construction, we see that $\text{Surger}(c)$ corresponds to the same cohomology class as c (integrating against $\text{Surger}(c)$ gives a map homotopic to φ), and so $\text{Surger}(c)$ represents the homology class x .

All of the choices made were unique up to isotopy, and so the 1-cycle $\text{Surger}(c)$ is a well-defined point of $\tilde{\mathcal{B}}(S)$.

In what follows, we will use the fact that, for any fixed 1-cycle c coming from a point of $\mathcal{B}(S)$, the function

$$d \mapsto \text{Surger}(tc + (1 - t)d)$$

is a continuous map from $\mathcal{B}(S) \times [0, 1]$ to $\tilde{\mathcal{B}}(S)$. For continuity, the key point is that the condition that some isotopy class is in the support of $\text{Surger}(tc + (1 - t)d)$ is an open condition.




FIGURE 5. Surgery on 1-cycles

The construction of $\text{Surger}(c)$ is more easily explained by a single picture; see Figure 5. The well-definedness is easy to see in this figure. What is not easy to see in the figure is that $\text{Surger}(c)$ is a 1-cycle. For instance, let α and β be two curves on the torus that intersect once, and let d be the 1-cycle $\epsilon\alpha + (1 - \epsilon)\beta$, where ϵ is irrational. If we perform the surgery indicated in Figure 5 at the intersection, the result is not a well-defined 1-cycle. The point is that d does not represent a rational homology class. In our construction of $\text{Surger}(c)$ it was essential that x was an integral class (a slight modification is needed in the rational case).

Draining 1-cycles. The result of the surgery defined above is always a simple positive 1-cycle representing x . However, it is not necessarily true that this 1-cycle represents a point of $\mathcal{B}(S)$; rather, this 1-cycle is in general a point in $\tilde{\mathcal{B}}(S)$. To address this issue, we now define a strong deformation retraction $\text{Drain} : \tilde{\mathcal{B}}(S) \rightarrow \mathcal{B}(S)$.

As a basic example, suppose that c is the formal sum of a 1-cycle representing a point of $\mathcal{B}(S)$ with a weighted oriented separating curve d . The cycle $\text{Drain}(c)$, in this case, is obtained by “draining” the coefficient of d to 0 (we think of the subsurface bounded by d as the drain).

We now return to the general case; let c be a 1-cycle representing a point of $\tilde{\mathcal{B}}(S)$. Let $\{R_i\}$ be the set of embedded subsurfaces of S that have boundary in c . Let $\{R_1, \dots, R_k\}$ be the subset of $\{R_i\}$ consisting of subsurfaces that give rise to one-sided relations. Suppose that the one-sided relation corresponding to R_1 is

$$[c_1] + \dots + [c_p] = 0.$$

For small enough ϵ , we obtain a new 1-cycle representing x by starting with c and decreasing each of the coefficients of c_1, \dots, c_p by ϵ .

We may simultaneously consider all of the k one-sided relations at once: again, we start with c , and we decrease the coefficient of c_i by $n\epsilon$, where $0 \leq n \leq k$ is the number of one-sided relations in which c_i appears.

We adjust weights in this way until some coefficient becomes 0. At this point, we repeat the process, starting with the new 1-cycle. When there are no more one-sided relations, we stop. The resulting 1-cycle is called $\text{Drain}(c)$. By Lemma 5.1, $\text{Drain}(c)$ represents a point of $\mathcal{B}(S)$. Note in particular, that, since no arbitrary choices were made, the 1-cycle $\text{Drain}(c)$ is well-defined.

The Drain map is a strong deformation retraction of $\tilde{\mathcal{B}}(S)$ onto $\mathcal{B}(S)$. Below, it will suffice to think of Drain as a continuous map $\tilde{\mathcal{B}}(S) \rightarrow \mathcal{B}(S)$.

The proof of contractibility. We choose an arbitrary basepoint v of $\mathcal{B}(S)$. We will now define a deformation retraction of $\mathcal{B}(S)$ to the point v .

Denote by c a 1-cycle representing v . The retraction $H : \mathcal{B}(S) \times [0, 1] \rightarrow \mathcal{B}(S)$ is given by

$$H(w, t) = \text{Drain}(\text{Surger}(tc + (1 - t)d)),$$

where w is an arbitrary point of $\mathcal{B}(S)$, and d is a representative 1-cycle.

The map H is a strong deformation retraction of $\mathcal{B}(S)$ onto the point v , and so $\mathcal{B}(S)$ is contractible.

6. STABILIZER DIMENSION

Recall from the Introduction that, to prove Theorems A and B, it suffices to show that the following inequalities hold for each cell σ of the complex of minimizing

cycles $\mathcal{B}(S_g)$:

$$\begin{aligned} \text{cd}(\text{Stab}_{\mathcal{I}(S_g)}(\sigma)) + \dim(\sigma) &\leq 3g - 5, \\ \text{cd}(\text{Stab}_{\mathcal{K}(S_g)}(\sigma)) + \dim(\sigma) &\leq 2g - 3. \end{aligned}$$

A cell σ of $\mathcal{B}(S_g)$ is associated to a multicurve M consisting of nonseparating curves. The stabilizer in $\mathcal{I}(S_g)$ of the cell σ is exactly the stabilizer of M . Since an element of $\mathcal{I}(S_g)$ cannot reverse the orientation of any simple closed curve, the orientation of M will play no role in this section.

As in Section 2, the dimension of the cell σ is the number $B(M)$, which only depends on M (without orientation).

In summary, the proofs of Theorems A and B are reduced to the following.

Proposition 6.1. *For $g \geq 2$ and M as above, $\text{cd}(\text{Stab}_{\mathcal{I}(S_g)}(M)) + B(M) \leq 3g - 5$.*

Proposition 6.2. *For $g \geq 2$ and M as above, $\text{cd}(\text{Stab}_{\mathcal{K}(S_g)}(M)) + B(M) \leq 2g - 3$.*

Before proving Propositions 6.1 and 6.2 in Sections 6.2 and 6.3, respectively, we introduce some preliminaries related to the Birman exact sequence.

6.1. The Birman exact sequence and dimension. In order to state the classical Birman exact sequence, we need one definition. The *pure mapping class group* of a surface S , denoted $\text{PMod}(S)$, is the subgroup of $\text{Mod}(S)$ consisting of elements that fix each puncture and boundary component of S individually.

Theorem 6.3 (Birman exact sequence). *Let S be a surface of negative Euler characteristic, and let S' be $S - p$. The following sequence is exact:*

$$1 \rightarrow \pi_1(S, p) \rightarrow \text{PMod}(S') \rightarrow \text{PMod}(S) \rightarrow 1.$$

The map $\text{PMod}(S') \rightarrow \text{PMod}(S)$ is obtained by “forgetting” the puncture p , and the image of an element of $\pi_1(S)$ in $\text{PMod}(S')$ is realized by “pushing” the puncture p around that element of $\pi_1(S)$; see [6, Section 4.1].

We will use the Birman exact sequence to describe the effect of punctures on the cohomological dimension of subgroups of the mapping class group. There are three statements: one for higher genus ($g \geq 2$), one for genus 1, and one for genus 0. We use $S_{g,p}$ to denote a surface of genus g with p punctures (and no boundary).

The three corollaries below follow, in a straightforward way, by inducting on the number of punctures, and applying the following facts; see [9, Chapter VIII, Proposition 2.4].

Fact 6.4 (Subadditivity of cohomological dimension). If we have a short exact sequence of groups

$$1 \rightarrow K \rightarrow G \rightarrow Q \rightarrow 1,$$

then $\text{cd}(G) \leq \text{cd}(K) + \text{cd}(Q)$.

Fact 6.5 (Monotonicity of cohomological dimension). If H is a subgroup of G , then $\text{cd}(H) \leq \text{cd}(G)$.

One also needs the basic facts that $\text{cd}(\pi_1(S_g)) = 2$ for $g \geq 1$, and $\text{cd}(\pi_1(S_{g,p})) = 1$ if $p > 0$, and $(g, p) \neq (0, 1)$.

We start with the case $g \geq 2$. For the application of Corollary 6.6 in Section 6.2, we will take $H = \mathcal{I}(S_g)$.

Corollary 6.6. *Let $g \geq 2$, and let G be a subgroup of $\text{PMod}(S_{g,p})$. Let $F : \text{PMod}(S_{g,p}) \rightarrow \text{Mod}(S_g)$ be the map induced by forgetting all of the punctures, and suppose $F(G)$ is contained in a subgroup H of $\text{Mod}(S_g)$. We have*

$$\text{cd}(G) \leq \text{cd}(H) + (p + 1).$$

Since the torus has Euler characteristic zero, we will instead apply Theorem 6.3 to the once-punctured torus. We have the following analogue of Corollary 6.6 for this scenario. In Section 6.2, we will apply Corollary 6.7 with $g = 1$, $k = 1$, and $H = \mathcal{I}(S_{1,1}) = 1$.

Corollary 6.7. *Let $g \geq 1$, let $p \geq k > 0$, and let $F : \text{PMod}(S_{g,p}) \rightarrow \text{Mod}(S_{g,k})$ be the map induced by forgetting $p - k$ of the punctures of $S_{g,p}$. If $G < \text{PMod}(S_{g,p})$ and $F(G) < H$, then*

$$\text{cd}(G) \leq \text{cd}(H) + (p - k).$$

In the genus 0 case, we only need the following.

Corollary 6.8. *For $p \geq 3$, we have $\text{cd}(\text{PMod}(S_{0,p})) \leq p - 3$.*

To prove the last corollary, one needs the fact that $\text{PMod}(S_{0,3}) = 1$.

In Section 6.3, we will need a specialized version of Corollary 6.6. In the statement, a *bounding pair map* is a product $T_c T_d^{-1}$, where c and d are disjoint, homologous nonseparating curves.

Lemma 6.9. *Let $g \geq 2$, and let G be a subgroup of $\text{PMod}(S_{g,p})$. Let $F_1 : \text{PMod}(S_{g,p}) \rightarrow \text{Mod}(S_{g,1})$ and $F : \text{PMod}(S_{g,p}) \rightarrow \text{Mod}(S_g)$ be the maps obtained by forgetting all but one of the punctures and all of the punctures, respectively. Suppose that $F_1(G)$ contains no power of any bounding pair map, and that $F(G)$ is contained in a subgroup H of $\text{Mod}(S_g)$. We have*

$$\text{cd}(G) \leq \text{cd}(H) + p.$$

Proof. As in the case of Corollary 6.6, we apply Theorem 6.3 inductively. However, to improve from $p + 1$ to p , we need an extra argument in the last step.

By the same straightforward applications of Fact 6.4 and Theorem 6.3 used for Corollary 6.6, we have that

$$\text{cd}(G) \leq \text{cd}(F_1(G)) + (p - 1).$$

It remains to show that

$$\text{cd}(F_1(G)) \leq \text{cd}(H) + 1.$$

We consider the exact sequence

$$1 \rightarrow A \rightarrow F_1(G) \rightarrow F(G) \rightarrow 1.$$

By the Birman exact sequence (Theorem 6.3), the kernel A is a subgroup of $\pi_1(S_g)$. We claim that A has infinite index in $\pi_1(S_g)$; indeed, under the map $\pi_1(S_{g,1}) \rightarrow \text{Mod}(S_{g,1})$ from Theorem 6.3, an element α of $\pi_1(S_{g,1})$ represented by a simple nonseparating loop maps to a bounding pair map (cf. [6, Figure 14]), and so by assumption no power of α is an element of A . It follows that A is a free group (the corresponding cover of S_g is noncompact). This implies that $\text{cd}(A) \leq 1$. It remains to apply Facts 6.4 and 6.5 (subadditivity and monotonicity). \square

6.2. Stabilizer dimensions for $\mathcal{I}(S_g)$. For the duration of this section, we fix some $g \geq 2$. Also, throughout the section, let M be a fixed multicurve in S_g consisting entirely of nonseparating curves. Suppose that $S_g - M$ has P components of positive genus R_1, \dots, R_P and Z components of genus zero R_{P+1}, \dots, R_{P+Z} . Say that R_i is homeomorphic to a surface of genus g_i with p_i punctures.

Let $G(M)$ be the free abelian group generated by the Dehn twists in the curves of M . It is a theorem of Vautaw that $G(M) \cap \mathcal{I}(S_g)$ is generated by bounding pair maps [46, Theorem 3.1]. Let BP be the number of curves of M minus the number of distinct homology classes represented by the curves of M . We will use the following consequence of Vautaw’s theorem.

Theorem 6.10 (Vautaw). *For M as above, we have $G(M) \cap \mathcal{I}(S_g) \cong \mathbb{Z}^{BP}$.*

We will apply Corollaries 6.6 and 6.7 in the case where H is the Torelli group. The work of Ivanov [26, Corollary 1.8] shows that the map from $\text{Stab}_{\mathcal{I}(S_g)}(M)$ to $\text{Mod}(S_g - M)$ has image in $\text{PMod}(S_g - M)$, and so for each i there is a map $\text{Stab}_{\mathcal{I}(S_g)}(M) \rightarrow \text{PMod}(R_i)$; in particular, the R_i are preserved. Let \hat{R}_i be the surface obtained from R_i by forgetting all of the punctures, and let \hat{R}'_i be the surface obtained from R_i by forgetting all of the punctures but one. For R equal to \hat{R}_i or \hat{R}'_i , there is an induced homomorphism $\text{PMod}(R_i) \rightarrow \text{Mod}(R)$.

Lemma 6.11. *Let R be either \hat{R}_i or \hat{R}'_i . The image of $\text{Stab}_{\mathcal{I}(S_g)}(M)$ in $\text{Mod}(R)$ is a subgroup of $\mathcal{I}(R)$.*

Proof. We start with the case of \hat{R}'_i . Let $\hat{f} \in \text{Mod}(\hat{R}'_i)$ be the image of $f \in \text{PMod}(R_i)$, and let \hat{c} be a simple closed curve in \hat{R}'_i . We wish to show that $\hat{f}(\hat{c})$ is homologous to \hat{c} . Let c be a curve in $R_i \subset S_g$ that maps to \hat{c} under the map $R_i \rightarrow \hat{R}'_i$. Since $f \in \mathcal{I}(S_g)$, the curves c and $f(c)$ are homologous, and so there is an immersed surface Σ in S_g that realizes this homology. By adding a multiple of the fundamental class of S_g if necessary, we may assume that Σ avoids any particular puncture of R_i . The surface \hat{R}'_i is obtained from R_i by gluing in points (at the punctures of R_i), and if we glue in points to Σ , we obtain a homology in \hat{R}'_i between \hat{c} and $\hat{f}(\hat{c})$.

The proof for the case of \hat{R}_i is the same as the above, but the step where Σ is modified is no longer necessary. □

Lemma 6.12. *Assume Proposition 6.1 for all genera between 2 and $g-1$, inclusive. In the notation of the preceding paragraphs, we have*

$$\text{cd}(\text{Stab}_{\mathcal{I}(S_g)}(M)) \leq \sum_{i=1}^P (3g_i + p_i - 4) + \sum_{i=P+1}^{P+Z} (p_i - 3) + BP.$$

Proof. Birman–Lubotzky–McCarthy proved that we have the following short exact sequence [7, Lemma 2.1]:

$$1 \rightarrow G(M) \rightarrow \text{Stab}_{\text{Mod}(S_g)}(M) \rightarrow \text{Mod}(S_g - M) \rightarrow 1.$$

Here, the group $\text{Stab}_{\text{Mod}(S_g)}(M)$ consists of those mapping classes that fix the isotopy class of M , and, as above, $G(M) \cong \mathbb{Z}^{|M|}$ is the group generated by the Dehn twists about the curves of M .

As in the paragraph before Lemma 6.11, there is a map $\text{Stab}_{\mathcal{I}(S_g)}(M) \rightarrow \text{PMod}(R_i)$. Applying Theorem 6.10, we arrive at the appropriate sequence for the Torelli group:

$$1 \rightarrow \mathbb{Z}^{BP} \rightarrow \text{Stab}_{\mathcal{I}(S_g)}(M) \rightarrow \text{PMod}(S_g - M).$$

By Fact 6.4 (subadditivity), Theorem 6.10, and the fact that $\text{cd}(\mathbb{Z}^{BP}) = BP$, it remains to show that the image of $\text{Stab}_{\mathcal{I}(S_g)}(M)$ in $\text{PMod}(S_g - M)$ has dimension at most

$$\sum_{i=1}^P (3g_i + p_i - 4) + \sum_{i=P+1}^{P+Z} (p_i - 3).$$

Let G_i be the image of $\text{Stab}_{\mathcal{I}(S_g)}(M)$ in each $\text{PMod}(R_i)$, and let G be the direct product of the G_i . Since the image of $\text{Stab}_{\mathcal{I}(S_g)}(M)$ in $\text{PMod}(S_g - M)$ is a subgroup of G , it suffices to show that if $g_i = 0$, then $\text{cd}(G_i) \leq p_i - 3$, and if $g_i > 0$, then $\text{cd}(G_i) \leq 3g_i + p_i - 4$ (apply Facts 6.5 and 6.4 (monotonicity and subadditivity)).

The genus zero case follows immediately from Fact 6.5 (monotonicity) and the inequality $\text{cd}(\text{PMod}(S_{0,p})) \leq p - 3$ (Corollary 6.8).

Now suppose R_i has genus 1. Let \hat{R}_i be the once-punctured torus obtained from R_i by forgetting all of the punctures except for one. By Lemma 6.11, the image of G_i in $\text{Mod}(\hat{R}_i)$ is a subgroup of $\mathcal{I}(\hat{R}_i)$. Also, it is a fact that $\mathcal{I}(S_{1,1}) = 1$. Thus, by Corollary 6.7, we have that $\text{cd}(G_i) \leq 0 + (p_i - 1)$. The latter happens to be equal to $3g_i + p_i - 4$.

The case $g_i \geq 2$ is similar. Here, we take \hat{R}_i to be the closed surface obtained from R_i by forgetting all of the punctures. Since $g_i < g$, we can apply Proposition 6.1, which tells us that $\text{cd}(\mathcal{I}(\hat{R}_i)) = 3g_i - 5$. Applying Lemma 6.11 and Corollary 6.6, we arrive at $\text{cd}(G_i) \leq (3g_i - 5) + (p_i + 1) = 3g_i + p_i - 4$, which completes the proof. \square

We now give another version of Lemma 6.12, obtained by algebraic manipulation of the upper bound.

Lemma 6.13. *With the same assumptions and notation as in Lemma 6.12, we have*

$$\text{cd}(\text{Stab}_{\mathcal{I}(S_g)}(M)) \leq 3g - 3 - P - |M| + BP.$$

Since the quantity BP is equal to $|M| - C$, where C is the number of distinct homology classes represented by the curves of M , the upper bound of the lemma is equivalent to $3g - 3 - P - C$. For the proof, recall that D is the dimension of the span of the classes of the curves of M in $H_1(S_g, \mathbb{R})$.

Proof. By Lemma 6.12, it suffices to show that

$$\sum_{i=1}^P (3g_i + p_i - 4) + \sum_{i=P+1}^{P+Z} (p_i - 3) + BP = 3g - 3 - P - |M| + BP.$$

We have

$$\begin{aligned}
 \sum_{i=1}^P (3g_i + p_i - 4) + \sum_{i=P+1}^{P+Z} (p_i - 3) + BP &= 3 \sum g_i + 2|M| - 4P - 3Z + BP \\
 &= 3(g - D) + 2|M| - 4P - 3Z + BP \\
 &= 3g - 3(|M| + 1) + 2|M| - P + BP \\
 &= 3g - 3 - P - |M| + BP.
 \end{aligned}$$

For the first equality, we used the fact that $\sum p_i = 2|M|$. Then, we applied the fact that $D = g - \sum g_i$. The third equality is a consequence of Lemma 2.1, which says that $D + P + Z = |M| + 1$, and the last equality is obtained by simple algebra. \square

The next lemma is a basic fact about multicurves made up of nonseparating curves.

Lemma 6.14. *For a multicurve M in S_g that consists entirely of nonseparating curves, we have $BP + 2 \leq D + P$.*

Proof. We start with an observation: for any multicurve in S_g with $g \geq 2$, we have $D + P \geq 2$ (recall that multicurves are nonempty).

We now proceed by induction on genus, with the base case $g = 2$. In that case, $BP = 0$ (there are no bounding pairs in S_2). Since $D + P \geq 2$, we are done in this case.

Now assume $g \geq 3$, and that the lemma is true for all genera between 2 and $g - 1$, inclusive. Let M be a multicurve in S_g . If there are no bounding pairs, then, as above, there is nothing to do. So assume $BP > 0$.

We can find an “innermost bounding pair” of M , by which we mean a bounding pair $\{c, d\} \subseteq M$ so that one component of $S_g - (c \cup d)$ contains no other bounding pairs of M . For example, in Figure 6, the bounding pairs $\{a_1, a_2\}$, $\{a_2, a_3\}$, and $\{c_1, c_2\}$ are innermost, whereas the bounding pairs $\{a_1, a_3\}$ and $\{b_1, b_2\}$ are not innermost.




FIGURE 6. A multicurve in S_6 with $BP = 4$, $D = 4$, and $P = 2$.

When there is at least one bounding pair, then there is an innermost one: if there is a bounding pair on one side of a given bounding pair, then the new bounding pair bounds a strictly smaller subsurface (as measured by the Euler characteristic), and so the process of finding increasingly innermore bounding pairs must terminate.

Let $\{c, d\}$ be an innermost bounding pair of M , and let R be a component of $S_g - (c \cup d)$ that contains no bounding pairs of M . We obtain a new surface S' from S_g by replacing R with an annulus. There is an induced multicurve M' on S' , after removing the image of either c or d (they become homotopic by construction).

We claim that the genus of S' lies in $[2, g - 1]$. Indeed, the closures of R and $S_g - R$ have two boundary components each, and so if either had genus zero, we would see that c and d were homotopic.

By induction, M' satisfies

$$BP' + 2 \leq D' + P',$$

where BP' , D' , and P' are defined for M' in the same way as BP , D , and P for M . Since $\{c, d\}$ was innermost in S , we have

$$BP' = BP - 1.$$

Also we have

$$D' + P' + 1 \leq D + P$$

because when we replace R with the annulus, we are either deleting a positive genus component of $S_g - M$, or we see that there is at least one nonseparating curve in R , which is automatically homologically independent of the other curves of M . In the latter case, replacing R by an annulus reduces D . Thus, we have

$$BP + 2 = BP' + 3 \leq D' + P' + 1 \leq D + P,$$

and the proof is complete. \square

We now prove Proposition 6.1, which completes the proof of Theorem A.

Proof of Proposition 6.1. We proceed by strong induction. Assume that Proposition 6.1 holds for genera between 2 and $g - 1$ inclusive (if $g = 2$ this is an empty assumption). We want to prove the following inequality:

$$\text{cd}(\text{Stab}_{\mathcal{I}(S_g)}(M)) + B(M) \leq 3g - 5.$$

We can replace $B(M)$ with $|M| - D$ (Lemma 2.1). Applying Lemma 6.13, it is enough to show that

$$(3g - 3 - P - |M| + BP) + (|M| - D) \leq 3g - 5.$$

But this inequality is equivalent to the inequality $BP + 2 \leq D + P$, and so an application of Lemma 6.14 completes the proof. \square

6.3. Stabilizer dimensions for $\mathcal{K}(S_g)$. Throughout this section we again fix some $g \geq 2$ and some multicurve M in S_g consisting entirely of nonseparating curves. As in Section 6.2, we label the positive genus components of $S_g - M$ by R_1, \dots, R_P , and the genus zero components by R_{P+1}, \dots, R_{P+Z} , and we say that R_i is a surface of genus g_i with p_i punctures. Recall that for a surface S , we have defined $\mathcal{K}(S)$ as the group generated by Dehn twists about separating curves in S .

We will require the analogue of Theorem 6.10 for $\mathcal{K}(S_g)$, proven as Theorem A.1(ii) in Appendix A.

Theorem 6.15. *For M as above, we have $G(M) \cap \mathcal{K}(S_g) = 1$.*

We will also need the analogue of Lemma 6.11 for the Johnson kernel. Let \hat{R}_i be the surface obtained from R_i by forgetting all of the punctures, and let \hat{R}'_i be the surface obtained by forgetting all of the punctures except for one. The next lemma is a direct consequence of the work of Church and Putman.

Lemma 6.16. *Let R be either \hat{R}_i or \hat{R}'_i . The image of $\text{Stab}_{\mathcal{K}(S_g)}(M)$ in $\text{Mod}(R)$ is a subgroup of $\mathcal{K}(R)$.*

Lemma 6.16 is stated explicitly by Church [11]. Putman has explained to us how to derive it from the Birman exact sequence for the Johnson kernel in his paper [39, Proposition 4.1].

Next, we have an analogue of Lemma 6.12.

Lemma 6.17. *Assume Proposition 6.2 for all genera between 2 and $g-1$, inclusive. In the notation of the preceding paragraph, we have*

$$\text{cd}(\text{Stab}_{\mathcal{K}(S_g)}(M)) \leq \sum_{i=1}^P (2g_i + p_i - 3) + \sum_{i=P+1}^{P+Z} (p_i - 3).$$

Proof. The proof is essentially the same as the proof of Lemma 6.12, with lemmas about $\mathcal{I}(S_g)$ replaced by lemmas about $\mathcal{K}(S_g)$.

As in the proof of Lemma 6.12, we cut S_g along M and obtain

$$1 \rightarrow G(M) \cap \mathcal{K}(S_g) \rightarrow \text{Stab}_{\mathcal{K}(S_g)}(M) \rightarrow \text{PMod}(S_g - M),$$

where $G(M)$ is the group generated by Dehn twists in the curves of M . By Theorem 6.15, the term $G(M) \cap \mathcal{K}(S_g)$ is trivial, and so $\text{Stab}_{\mathcal{K}(S_g)}(M)$ is isomorphic to its image in $\text{PMod}(S_g - M)$. So we must show that this image has cohomological dimension bounded above by

$$\sum_{i=1}^P (2g_i + p_i - 3) + \sum_{i=P+1}^{P+Z} (p_i - 3).$$

As in the proof of Lemma 6.12, we denote by G_i the image of $\text{Stab}_{\mathcal{K}(S_g)}(M)$ in $\text{PMod}(R_i)$.

As before, it suffices to show that if $g_i = 0$, then $\text{cd}(G_i) \leq p_i - 3$, and if $g_i > 0$, then $\text{cd}(G_i) \leq 2g_i + p_i - 3$.

As in Lemma 6.12, the genus zero case follows from Corollary 6.8. For $g_i = 1$, note that $2g_i + p_i - 3 = 3g_i + p_i - 4 = p_i - 1$, and so, by Fact 6.5 (monotonicity), the upper bound of $p_i - 1$ obtained from Corollary 6.7 in the proof of Lemma 6.12 is sufficient.

It remains to deal with the case $g_i \geq 2$. Let \hat{R}_i be the closed surface obtained from R_i by forgetting all of the punctures, and let \hat{R}'_i be the surface obtained from R_i by forgetting all punctures but one. Let \hat{G}_i and \hat{G}'_i be the images of G_i in $\text{Mod}(\hat{G}_i)$ and $\text{Mod}(\hat{G}'_i)$, respectively. By Lemma 6.16, we have $\hat{G}_i < \mathcal{K}(\hat{R}_i)$ and $\hat{G}'_i < \mathcal{K}(\hat{R}'_i)$. By Theorem 6.15, $\mathcal{K}(\hat{R}'_i) \cong \mathcal{K}(S_{g_i,1})$ contains no nontrivial power of any bounding pair map (see also [27, Lemmas 4A and 4B]). Thus, Lemma 6.9 applies, and we have $\text{cd}(G_i) \leq \text{cd}(\mathcal{K}(\hat{R}_i)) + p_i$. By induction, $\text{cd}(\mathcal{K}(\hat{R}_i)) \leq 2g_i - 3$, and this completes the proof. \square

We can simplify the upper bound of Lemma 6.17 as follows. The proof is analogous to the argument for Lemma 6.13 and is left to the reader.

Lemma 6.18. *With the assumptions and notation of Lemma 6.17, we have*

$$\text{cd}(\text{Stab}_{\mathcal{K}(S_g)}(M)) \leq 2g - 3 + D - |M|.$$

We are now ready to prove the upper bound for Theorem B.

Proof of Proposition 6.2. We proceed as in the proof of Proposition 6.1. Assume by induction that Proposition 6.2 is true for all genera between 2 and $g - 1$, inclusive. The goal is to show the following:

$$\text{cd}(\text{Stab}_{\mathcal{K}(S_g)}(M)) + B(M) \leq 2g - 3.$$

By Lemma 6.18 and Lemma 2.1 it suffices to show that

$$(2g - 3 + D - |M|) + (|M| - D) \leq 2g - 3.$$

But this inequality simplifies to $0 \leq 0$. □

Remark. We point out that in the case of $\mathcal{K}(S_g)$ we did not need an analogue of Lemma 6.14. In this sense, the upper bound for Theorem B is slightly easier than that for Theorem A. Also recall that the lower bound was also easier for Theorem B.

7. THE MESS DESCRIPTION OF $\mathcal{I}(S_2)$

We now turn to the proof of Theorem D. As discussed in the introduction, Theorem A implies that $\mathcal{I}(S_2)$ is a free group. However, the statement of Theorem D is much stronger, and the infinite generation is needed for the base case of Theorem C.

The proof breaks up into three parts: we first analyze the stabilizer in $\mathcal{I}(S_2)$ of a pair of nonseparating curves, then we analyze the stabilizer of a single nonseparating curve, and finally we piece these together to get the desired description of $\mathcal{I}(S_2)$. We can think of these steps as the genus 0, genus 1, and genus 2 versions of the Mess theorem.

Genus 0. We start by investigating the stabilizer in $\mathcal{I}(S_2)$ of a vertex of $\mathcal{B}(S_2)$ corresponding to a pair of disjoint nonseparating curves. Since the complement in S_2 has genus zero, we think of this as a genus zero version of Theorem D.

In the statement of the lemma, we say that a symplectic splitting of $H_1(S_2, \mathbb{Z})$ is *compatible* with a pair of nonseparating curves $\{a, b\}$ (or the associated homology classes) if the homology class $[a]$ lies in one subspace determined by the splitting, and $[b]$ lies in the other.

Lemma 7.1. *Let M be the multicurve consisting of two nonseparating curves a and b in S_2 . The group $\text{Stab}_{\mathcal{I}(S_2)}(M)$ is an infinitely generated free group, with one Dehn twist generator for each symplectic splitting compatible with M .*

Proof. We cut S_2 along a and b and obtain a sphere S' with 4 punctures. Let X be the flag complex with a vertex for each isotopy class of curves in S' that comes from a separating curve in S_2 , and an edge for each pair of isotopy classes with geometric intersection number 4.

The complex X is contractible: if we think of the vertices as isotopy classes of arcs in S' connecting the two punctures coming from a , then edges correspond to disjointness of arcs, and the required statement is given by Harer [18, Theorem 1.6] (the argument of [22] also applies, with minor modifications). We also see from this point of view that X is a graph (so it is a tree): any three arcs in S' that represent distinct separating curves must intersect.

The group $\text{Stab}_{\mathcal{I}(S_2)}(M)$ acts on X . The quotient has one vertex for each homology splitting of $H_1(S_2, \mathbb{Z})$ compatible with M ; in other words, if two separating curves c and c' are disjoint from M and induce the same homology splitting, then there is an element of $\mathcal{I}(S_2)$ that fixes a and b and takes c to c' . To construct such a mapping class, cut one copy of S along c and another copy of S along c' , choose the unique maps on the punctured tori that act trivially on homology, and then choose any element of $\text{Mod}(S_2)$ that restricts to the chosen maps on the punctured tori (cf. [28, Theorem 1A] and [40, Lemma A.3]).

Our goal is to show that the quotient complex is contractible (a tree). Since vertex stabilizers are the Dehn twists about the corresponding separating curves in S_2 , and edge stabilizers are trivial (Theorem 6.10, plus the fact that any two nonisotopic separating curves in S_2 must intersect), the lemma will then follow from the classical Bass–Serre theory of group actions on trees [42, Theorem 13].

Let $([a], [a'], [b], [b'])$ be a symplectic basis for $H_1(S_2, \mathbb{Z})$, where a and b are the curves in the statement of the lemma (take the algebraic intersection of $[a]$ and $[a']$, say, to be 1). It follows from elementary linear algebra that the symplectic splittings of $H_1(S_2, \mathbb{Z})$ compatible with M are in bijection with the cosets $([a'] + k[b]) + \langle [a] \rangle$ and hence are indexed by \mathbb{Z} .

We now claim that two vertices form an edge in the quotient complex if and only if they are adjacent in \mathbb{Z} . It will follow that the quotient complex is isomorphic to the standard triangulation of \mathbb{R} .




FIGURE 7. The arcs α_0 and $T_c(\alpha_0)$ from the proof of Lemma 7.1.

If two vertices are adjacent in \mathbb{Z} , then one can draw disjoint arcs in S' corresponding to those vertices. This is proven by drawing the picture: if α_0 is an arc in S' corresponding to the integer 0, then $T_c^k(\alpha_0)$ is an arc corresponding to k , where T_c is the Dehn twist about any fixed curve in S' that comes from a curve in S_2 , represents the class $[a] + [b]$, and intersects α_0 in one point; see Figure 7 for an illustration. Since α_0 and $T_c(\alpha_0)$ are disjoint (by inspection), it follows that $T_c^k(\alpha_0)$ and $T_c^{k+1}(\alpha_0)$ are disjoint, and so the vertices corresponding to k and $k + 1$ are connected by an edge in the quotient.

If two vertices labelled by $k, k' \in \mathbb{Z}$ are connected in the quotient, this means that there are corresponding arcs in S' that are disjoint, and hence corresponding curves in S_2 that are disjoint. As above, these curves represent primitive homology classes $[a'] + k[b] + j[a]$ and $[a'] + k'[b] + j'[a]$. Since the curves are disjoint, we have $j = j'$ (otherwise, the algebraic intersection is nonzero). The curves in S_2 span a maximal isotropic subspace (i.e., a maximal subspace where the intersection form is trivial); in particular, they span the subspace spanned by $[a']$ and $[b]$. Hence, the

determinant of the pair of vectors $[a'] + k[b]$ and $[a'] + k'[b]$ must be ± 1 , and we see that $k = k' \pm 1$, which is what we wanted to show. \square

Genus 1. We now need to understand the stabilizer in $\mathcal{I}(S_2)$ of a vertex of $\mathcal{B}(S_2)$ corresponding to a single nonseparating curve c in the class x . This is the genus 1 version of Theorem D. We say that a homology splitting is *compatible with x* if x lies in one component of the splitting.

Lemma 7.2. *The group $\text{Stab}_{\mathcal{I}(S_2)}(c)$ is an infinitely generated free group, with one Dehn twist generator for each symplectic splitting compatible with x .*

Proof. We cut S_2 along c and obtain a twice-punctured torus S' . To find a generating set for $\text{Stab}_{\mathcal{I}(S_2)}(c)$, we consider the action on a complex of minimizing cycles for S' . To this end, we fix a homology class $y \in H_1(S_2, \mathbb{Z})$ coming from a nonseparating curve d in S_2 that is disjoint (and isotopically distinct) from c . The classes x and y correspond to classes in $H_1(S', \mathbb{Z})$, which we also call x and y .

The choice of the class y gives rise to a complex of minimizing cycles $\mathcal{B}(S')$, which is contractible (by Section 5 or by modifying the proof in Sections 3 and 4).

Given any two disjoint, isotopically distinct, nonperipheral nonseparating curves in S' , their homology classes necessarily differ by $\pm x$. It follows that an integral basic cycle for y consists of a single curve in the homology class $y + kx$ (along with $-kc$) and that $\mathcal{B}(S')$ is a tree. We also deduce that the vertices of the quotient are indexed by \mathbb{Z} , and that if two vertices in the quotient are connected by an edge, then they are adjacent in \mathbb{Z} . Also, if two vertices are adjacent in \mathbb{Z} , then we can realize them by disjoint curves. Thus the quotient of $\mathcal{B}(S')$ by $\text{Stab}_{\mathcal{I}(S_2)}(c)$ is isomorphic to the standard triangulation of \mathbb{R} .

As in Lemma 7.1, $\text{Stab}_{\mathcal{I}(S_2)}(c)$ is freely generated by the vertex stabilizers (we take one copy for each vertex of the quotient). By Lemma 7.1, each vertex stabilizer is an infinitely generated free group, with one Dehn twist generator for each homology splitting compatible with x and $y + kx$.

It remains to establish the bijection between generators of $\text{Stab}_{\mathcal{I}(S_2)}(c)$ and homology splittings compatible with x . By Lemma 7.1, we get, for each k , one generator for each splitting compatible with x and $y + kx$. This accounts for all splittings compatible with x since any splitting compatible with x splits y into $y + kx$ and $-kx$. \square

Genus 2. We finally use Lemmas 7.1 and 7.2 to obtain the full version of the Mess theorem.

Proof of Theorem D. By Theorem E, the complex of minimizing cycles $\mathcal{B}(S_2)$ is contractible. We also have that it is one dimensional by Lemma 2.1, and so it is in fact a tree. As in the previous lemmas, we need to argue that the quotient $\mathcal{B}(S_2)/\mathcal{I}(S_2)$ is contractible, and then do some bookkeeping to check that there is exactly one generator for each homology splitting.

The vertices of $\mathcal{B}(S_2)/\mathcal{I}(S_2)$ are in one-to-one correspondence with the homology classes of basic cycles for x ; that is, if two nonseparating curves, or two pairs of nonseparating curves, in S_2 represent the same homology classes, then there is an element of $\mathcal{I}(S_2)$ taking one to the other [40, Lemma A.3]. There is a single distinguished vertex corresponding to the homology class x , and infinitely many vertices corresponding to pairs of homology classes (a basic cycle in S_2 has one or two curves).

In order to show that $\mathcal{B}(S_2)/\mathcal{I}(S_2)$ is contractible, we will assign a “weight” to each vertex and show that each vertex is connected by an edge to exactly one vertex of smaller weight. We declare the weight of the distinguished vertex corresponding to a single curve in the class x to be 1. For a vertex given by a basic cycle $pa + qb$ with $p, q > 0$, we define the weight to be $p + q$.

Let v be the vertex corresponding to the positive cycle $pa + qb$. There are exactly two classes that can be represented disjointly from the multicurve $\{a, b\}$, namely, $[a] + [b]$ and $[a] - [b]$. Hence, there are four potentially adjacent vertices, given by the pairings of $[a]$ and $[b]$ with these two new classes.

The class $[a] - [b]$ cannot pair with either $[a]$ or $[b]$ to give a vertex of smaller weight. Indeed, we have $(p + q)[a] + q([b] - [a]) = x$ and $(p + q)[b] + p([a] - [b]) = x$, and so in either case the weight is larger than $p + q$.

The weights of the other two nearby vertices are both smaller than $p + q$, since we have $(q - p)[b] + p([a] + [b]) = x$ and $(p - q)[a] + q([a] + [b]) = x$. As $q - p$ is negative, it follows from Lemma 3.9 that v is not connected by an edge to the vertex corresponding to the first equation. Since $p - q$ is positive, it follows from Lemma 3.13 that v is connected to the other vertex. This completes the proof of contractibility of $\mathcal{B}(S_2)/\mathcal{I}(S_2)$.

Now for the bookkeeping. The stabilizer of any lift of the distinguished vertex is an infinitely generated free group generated by one Dehn twist generator for each homology splitting compatible with x by Lemma 7.2. The stabilizer of a lift of a vertex corresponding to two homology classes is an infinitely generated free group generated by one Dehn twist for each splitting compatible with that pair by Lemma 7.1. There is no overlap between these two sets of generators, since they correspond to the splittings that are compatible with x and that are not compatible with x , respectively. The theorem follows. \square




FIGURE 8. A piece of $\mathcal{B}(S_2)/\mathcal{I}(S_2)$ corresponding to a particular Lagrangian subspace of $H_1(S_2, \mathbb{Z})$. The Farey graph is drawn lightly for reference.

A picture of the quotient. It is more illuminating to draw a diagram of the quotient $\mathcal{B}(S_2)/\mathcal{I}(S_2)$. It is naturally subdivided into pieces, corresponding to different 2-dimensional isotropic subspaces of $H_1(S_2, \mathbb{Z})$ that contain x . Each such subspace is group isomorphic to \mathbb{Z}^2 . The set of bases for \mathbb{Z}^2 is depicted via the Farey graph: vertices correspond to (unsigned) primitive vectors, and edges correspond to bases. There is a distinguished vertex x , and the rest of the component of $\mathcal{B}(S_2)/\mathcal{I}(S_2)$ corresponding to this subspace is the tree shown in Figure 8. The entire complex $\mathcal{B}(S_2)/\mathcal{I}(S_2)$ is obtained by gluing infinitely many of these trees along their distinguished vertices.

8. INFINITE GENERATION OF HOMOLOGY

The goal of this section is to prove Theorem C, which states that, for $g \geq 2$, the group $H_{3g-5}(\mathcal{I}(S_g), \mathbb{Z})$ is infinitely generated. The basic idea is to use induction, with base case Theorem D, and the spectral sequences associated to the appropriate Birman exact sequences.

Throughout this section, we will be forced to consider homology with various local and global coefficients. When the coefficient ring is \mathbb{Z} , it will be convenient to omit this from the notation.

8.1. Spectral sequences. We recall some facts about spectral sequences; for an introduction to the theory, see [20], [16], or [9].

Hochschild–Serre spectral sequence. Given a short exact sequence of groups

$$1 \rightarrow K \rightarrow G \rightarrow Q \rightarrow 1,$$

there is an associated Hochschild–Serre spectral sequence. The second page of this spectral sequence is the two-dimensional array of homology groups

$$E_{p,q}^2 = \begin{cases} H_p(Q, H_q(K)) & p, q \geq 0, \\ 0 & \text{otherwise,} \end{cases}$$

where the coefficients $H_q(K)$ are considered to be local: Q acts on K by conjugation. The main difficulties in calculation come from this action.

We remark that Fact 6.4 is an easy consequence of the Hochschild–Serre spectral sequence.

Fact 8.1. Consider the Hochschild–Serre spectral sequence associated to a short exact sequence

$$1 \rightarrow K \rightarrow G \rightarrow Q \rightarrow 1$$

and assume $\text{cd}(Q)$ and $\text{cd}(K)$ are finite. For $p = \text{cd}(Q)$ and $q = \text{cd}(K)$, we have

$$H_{p+q}(G) \cong E_{p,q}^2 = H_p(Q, H_q(K)).$$

Cartan–Leray spectral sequence. For a group G acting cellularly and without rotations on a contractible complex X , we can associate the Cartan–Leray spectral sequence (to act *without rotations* means that, if a cell is fixed by some element, then the cell is fixed pointwise by that element). The first page of this spectral sequence is given by

$$E_{p,q}^1 = \begin{cases} \bigoplus_{\sigma \in X_p} H_q(G_\sigma) & p, q \geq 0, \\ 0 & \text{otherwise,} \end{cases}$$

where X_p is the set of p -cells of X/G , and G_σ is the stabilizer in G of a lift of the cell σ .

We only need the following basic fact; see [9, Section VII.7] and [16, Section 18].

Fact 8.2. Consider the Cartan–Leray spectral sequence associated to the action of the group G on the contractible complex X . Suppose that $E_{p,q}^1 = 0$ for $p + q > k$. Then $E_{0,k}^1 = \bigoplus H_k(G_v)$ injects as a subgroup of $H_k(G)$ (the sum is over all vertices v of X/G).

We remark that Quillen’s upper bound on the cohomological dimension of a group given in the introduction is an immediate consequence of the Cartan–Leray spectral sequence; this is Exercise 4 in Section VIII.2 of [9].

8.2. Torelli Birman exact sequences. In this section, we give two versions of the Birman exact sequence that are particular to the Torelli group: the first, due to Putman, allows us to pass from two punctures to one puncture; the second, due to Johnson, enables us to go from there to the closed case. For each sequence, we analyze the action of the quotient on the kernel.

Let v be a vertex of $\mathcal{B}(S_g)$ corresponding to a single curve c (the image of v in $\mathcal{B}(S_g)/\mathcal{I}(S_g)$ is the distinguished vertex). We can think of the stabilizer in $\mathcal{I}(S_g)$ of v (equivalently, of c) as a subgroup of $\text{Mod}(S_{g-1,2})$, and so, as in the Birman exact sequence (Theorem 6.3), there is a forgetful map

$$\text{Stab}_{\mathcal{I}(S_g)}(v) \rightarrow \text{Mod}(S_{g-1,1})$$

whose kernel is a subgroup of $\pi_1(S_{g-1,1})$.

Lemma 8.3. *The image of $\text{Stab}_{\mathcal{I}(S_g)}(v)$ in $\text{Mod}(S_{g-1,1})$ is a subgroup of $\mathcal{I}(S_{g-1,1})$. Moreover, the sequence*

$$1 \rightarrow K \rightarrow \text{Stab}_{\mathcal{I}(S_g)}(v) \rightarrow \mathcal{I}(S_{g-1,1}) \rightarrow 1,$$

where K is the commutator subgroup of $\pi_1(S_{g-1,1})$, is exact.

Proof. The first statement of the lemma is proven in the same way as Lemma 6.11. The surjectivity of the map $\text{Stab}_{\mathcal{I}(S_g)}(v) \rightarrow \mathcal{I}(S_{g-1,1})$ can be seen as follows: embed $S_{g-1,1}$ in $S_{g-1,2}$ so that $S_{g-1,2} - S_{g-1,1}$ is a twice-punctured disk; given an element of $\mathcal{I}(S_{g-1,1})$, we can choose a representative that extends by the identity, and this gives an element of $\text{Stab}_{\mathcal{I}(S_g)}(v)$ (the choice is not canonical). The description of K is due to Putman [40, Theorem 4.1] (Putman uses surfaces with boundary instead of with punctures, but only inconsequential changes to his proof are required). \square

We also have the following short exact sequence, due to Johnson [29].

Lemma 8.4. *For $g \geq 2$, we have*

$$1 \rightarrow \pi_1(S_g) \rightarrow \mathcal{I}(S_{g,1}) \rightarrow \mathcal{I}(S_g) \rightarrow 1.$$

As indicated in Section 8.1, we can apply the Hochschild–Serre spectral sequence to understand the groups in Lemmas 8.3 and 8.4. Since coefficients are local, we will need to understand the actions of the quotient groups on the kernels of the two sequences.

In the next lemma, we consider the action (by conjugation) of $\mathcal{I}(S_g)$ on $H_i(\pi_1(S_g)) \cong H_i(S_g)$ coming from the sequence in Lemma 8.4. This is the restriction of the usual action of $\text{Mod}(S_g)$ on $H_1(S_g)$.

Lemma 8.5. *The action of $\mathcal{I}(S_g)$ on $H_i(S_g)$ is trivial for $i = 0, 1, 2$.*

Proof. The case $i = 0$ is trivial. When $i = 1$, the statement of the lemma is the definition of $\mathcal{I}(S_g)$. For $i = 2$, the lemma is equivalent to the fact that elements of $\mathcal{I}(S_g)$ preserve the orientation of S_g . \square

We now consider the action of $\mathcal{I}(S_{g-1,1})$ on $H_1(K)$ coming from Lemma 8.3. Unlike the previous case, the action here is nontrivial (for instance, take an element of K that is the commutator of two simple loops that intersect only at the basepoint, and consider the action by a Dehn twist about a separating curve that is disjoint from the first loop and intersects the other one twice). However, we will be able to find a subgroup of $H_1(K)$ on which $\mathcal{I}(S_{g-1,1})$ acts trivially (Lemma 8.7 below). In the proof of Theorem C, this will allow us to use global coefficients.

First, we have a technical lemma about commutators.

Lemma 8.6. *Let $F = \langle x_1, \dots, x_{2n} \rangle$ be the free group generated by the x_i , let K be the commutator subgroup of F , and let $h \in K$ be the product of commutators $[x_1, x_2] \cdots [x_{2n-1}, x_{2n}]$. The class of h in $H_1(K)$ is nontrivial.*

Proof. The group $H_1(K)$ can be thought of as the first homology of the 1-skeleton of the standard cubing of \mathbb{R}^{2n} . It is clear now that h represents a nontrivial class. \square

As above, we consider the action by conjugation of $\mathcal{I}(S_{g-1,1})$ on $H_1(K)$ coming from Lemma 8.3. We now give an invariant submodule of $H_1(K)$.

Lemma 8.7. *There is a nontrivial subgroup of $H_1(K)$ on which $\mathcal{I}(S_{g-1,1})$ acts trivially.*

Proof. Let $\alpha_1, \beta_1, \dots, \alpha_{g-1}, \beta_{g-1}$ be the usual generators for $\pi_1(S_{g-1,1})$, and let δ be the following element of $\pi_1(S_{g-1,1})$:

$$\delta = [\alpha_1, \beta_1] \cdots [\alpha_{g-1}, \beta_{g-1}].$$

Clearly, $\delta \in K$. We will show that the class of δ in $H_1(K)$ is fixed by $\mathcal{I}(S_{g-1,1})$. By Lemma 8.6, this is a nontrivial element of $H_1(K)$, and so the module generated by this element is the desired submodule of $H_1(K)$.

Thinking of $\text{Stab}_{\mathcal{I}(S_g)}(v)$ as a subgroup of $\text{Mod}(S_{g-1,2})$, the image of δ under the map $K \rightarrow \text{Stab}_{\mathcal{I}(S_g)}(v)$ in Lemma 8.3 is the Dehn twist T_d about a curve d that bounds a twice-punctured disk in $S_{g-1,2}$ (we think of δ as the simple loop in this disk that is based at one puncture and surrounds the other).

To prove the lemma, it suffices to show that, given an arbitrary element f of $\mathcal{I}(S_{g-1,1})$, there is a lift F of f to $\text{Stab}_{\mathcal{I}(S_g)}(v)$ (thought of as a subgroup of $\text{Mod}(S_{g-1,2})$) so that $FT_dF^{-1} = T_{F(d)}$ is conjugate in K to T_d . In fact, we will be able to construct the lift F so that $F(d) = d$, and so $T_{F(d)}$ is actually equal to T_d .

Let $S_{g-1,1} \rightarrow S_{g-1,2}$ be a fixed embedding with the property that $S_{g-1,2} - S_{g-1,1}$ is the twice-punctured disk bounded by d . Given $f \in \mathcal{I}(S_{g-1,1})$, we take any representative homeomorphism that extends by the identity on all of $S_{g-1,2}$, and call the resulting mapping class F . By construction, $F(d) = d$. It is clear that F maps to f under the map $\text{Mod}(S_{g-1,2}) \rightarrow \text{Mod}(S_{g-1,1})$ that forgets one puncture of $S_{g-1,2}$. It remains to check that F is an element of $\text{Stab}_{\mathcal{I}(S_g)}(v)$.

We embed $S_{g-1,2}$ into S_g so that the punctures correspond to homotopic non-separating curves (that is, “uncut”). If the isotopy class corresponding to the punctures is c , we can extend F to an element of $\text{Mod}(S_g)$ that is well-defined up

to Dehn twists about c . Since the support of F is on a surface of genus $g - 1$ and one puncture, and F acts trivially on the homology of this subsurface, we can choose a lift of F to $\text{Mod}(S_g)$ that lies in $\mathcal{I}(S_g)$. This completes the proof. \square

8.3. Proof of infinite generation of top homology. This section contains the proof of Theorem C, which states that $H_{3g-5}(\mathcal{I}(S_g), \mathbb{Z})$ is infinitely generated.

Lemma 8.8. *If $H_{3g-5}(\mathcal{I}(S_g))$ is infinitely generated, then $H_{3g-3}(\mathcal{I}(S_{g,1}))$ is infinitely generated.*

Proof. We know that $\text{cd}(\pi_1(S_g)) = 2$. By Theorem A, $\text{cd}(\mathcal{I}(S_g)) = 3g - 5$. Applying Fact 8.1 to the short exact sequence of Lemma 8.4, and using the identification $H_2(\pi_1(S_g)) = H_2(S_g) \cong \mathbb{Z}$, we find

$$H_{3g-3}(\mathcal{I}(S_{g,1})) \cong H_{3g-5}(\mathcal{I}(S_g), \mathbb{Z}).$$

By Lemma 8.5, the coefficients in the last homology group are global. Therefore, by Theorem C, we have that $H_{3g-3}(\mathcal{I}(S_{g,1}))$ is infinitely generated. \square

We record the following corollary of Lemma 8.8 (use Lemma 8.4 and Fact 6.4).

Corollary 8.9. *For $g \geq 2$, we have $\text{cd}(\mathcal{I}(S_{g,1})) = 3g - 3$.*

Lemma 8.10. *If v is any vertex of $\mathcal{B}(S_g)$, then $H_{3g-5}(\text{Stab}_{\mathcal{I}(S_g)}(v))$ injects into $H_{3g-5}(\mathcal{I}(S_g))$.*

Proof. We consider the Cartan–Leray spectral sequence associated to the action of $\mathcal{I}(S_g)$ on $\mathcal{B}(S_g)$. By Proposition 6.1, we know that $E_{p,q}^1 = 0$ whenever $p + q > 3g - 5$. By Fact 8.2, the group $E_{0,3g-5}^1 = \bigoplus H_{3g-5}(\text{Stab}_{\mathcal{I}(S_g)}(v))$, where the sum is over a set of representatives of the vertices of $\mathcal{B}(S_g)/\mathcal{I}(S_g)$, injects into $H_{3g-5}(\mathcal{I}(S_g))$. The lemma follows. \square

We are now ready to prove the theorem.

Proof of Theorem C. We proceed by induction on the genus g . Theorem D tells us that $\mathcal{I}(S_2)$ is an infinitely generated free group, and so $H_1(\mathcal{I}(S_2))$ is infinitely generated. Let $g \geq 3$, and assume the theorem is true for $g - 1$; that is, we assume that $H_{3g-8}(\mathcal{I}(S_{g-1}))$ is infinitely generated.

Let v be a vertex of $\mathcal{B}(S_g)$ represented by a single curve in the class x . By Lemma 8.10, it suffices to show that $H_{3g-5}(\text{Stab}_{\mathcal{I}(S_g)}(v))$ is infinitely generated.

Consider the short exact sequence in Lemma 8.3. Since K is free, we have $\text{cd}(K) = 1$, and Corollary 8.9 gives $\text{cd}(\mathcal{I}(S_{g-1,1})) = 3g - 6$. Applying Fact 8.1 to the sequence in Lemma 8.3, we obtain

$$H_{3g-5}(\text{Stab}_{\mathcal{I}(S_g)}(v)) \cong H_{3g-6}(\mathcal{I}(S_{g-1,1}), H_1(K)).$$

Thus, the theorem is reduced to showing that $H_{3g-6}(\mathcal{I}(S_{g-1,1}), H_1(K))$ is infinitely generated. By Lemma 8.7, there is a nontrivial submodule M of $H_1(K)$ on which $\mathcal{I}(S_{g-1,1})$ acts trivially. Since M is torsion free (it is a subgroup of the first homology of a free group), the universal coefficient theorem gives

$$H_{3g-6}(\mathcal{I}(S_{g-1,1}), M) \cong H_{3g-6}(\mathcal{I}(S_{g-1,1})) \otimes M.$$

By the inductive hypothesis and Lemma 8.8, the latter is infinitely generated.

Now, the short exact sequence of modules

$$1 \rightarrow M \rightarrow H_1(K) \rightarrow H_1(K)/M \rightarrow 1$$

induces a long exact sequence of homology groups, in which we find the following:

$$\begin{aligned} \cdots \rightarrow H_{3g-5}(\mathcal{I}(S_{g-1,1}), H_1(K)/M) &\rightarrow H_{3g-6}(\mathcal{I}(S_{g-1,1}), M) \\ &\rightarrow H_{3g-6}(\mathcal{I}(S_{g-1,1}), H_1(K)) \rightarrow \cdots \end{aligned}$$

By Corollary 8.9, the first term shown is trivial. Thus, the infinitely generated group $H_{3g-6}(\mathcal{I}(S_{g-1,1}), M)$ injects into $H_{3g-6}(\mathcal{I}(S_{g-1,1}), H_1(K))$, and the theorem follows. \square

In general, the vertices of $\mathcal{B}(S_g)$ correspond to basic cycles for x supported on k curves, with $1 \leq k \leq g$. When $k < g$, the cohomological dimension of the vertex stabilizer is $3g - 4 - k$. By an argument similar to that of Theorem C, the top dimensional integral homology of a vertex stabilizer ($k < g$) is infinitely generated. For $k > 1$, there are infinitely many vertices of that type. Therefore, to prove that $H_{3g-4-k}(\mathcal{I}(S_g), \mathbb{Z})$ is infinitely generated, one would only have to show that the quotient of each vertex group by its incoming edge groups is nontrivial.

Question 8.11. Is it true that $H_{3g-4-k}(\mathcal{I}(S_g), \mathbb{Z})$ is infinitely generated for $2 \leq k \leq g - 1$?

9. TORELLI MESS SUBGROUPS

We now briefly describe Mess’s proof of Theorem 1.1. We do this for completeness, and also because Mess’s argument is not published in the exact form that gives Theorem 1.1. His original paper is a preprint, available from Institut des Hautes Études Scientifiques [36].

The basic idea is to inductively define subgroups Γ_g of $\mathcal{I}(S_g)$ with the property that $\text{cd}(\Gamma_g) = 3g - 5$. The theorem then follows from Fact 6.5 (monotonicity).

We employ the theory of Poincaré duality groups. A group Γ is a Poincaré duality group if the trivial $\mathbb{Z}\Gamma$ -module \mathbb{Z} admits a finite length projective resolution by finitely generated $\mathbb{Z}\Gamma$ -modules, and $H^i(\Gamma, \mathbb{Z}\Gamma)$ is trivial in all but one dimension, where it is group isomorphic to \mathbb{Z} ; see [31] and [4] for background. Fundamental groups of closed aspherical manifolds are Poincaré duality groups, for example.

We will need the following fact [31, Theorem 3].

Theorem 9.1. *Given a short exact sequence of groups*

$$1 \rightarrow K \rightarrow G \rightarrow Q \rightarrow 1,$$

where K and Q are Poincaré duality groups, it follows that G is a Poincaré duality group and $\text{cd}(G) = \text{cd}(K) + \text{cd}(Q)$.

To start the inductive process, we define Γ_2 to be the group generated by a single Dehn twist about a separating curve in S_2 . Since $\Gamma_2 \cong \pi_1(S^1) \cong \mathbb{Z}$, we have that Γ_2 is a Poincaré duality group. Of course, $\text{cd}(\Gamma_2) = 1$.

We now assume that Γ_g is given, and that it is a Poincaré duality group with cohomological dimension $3g - 5$. We will define Γ_{g+1} in two steps.

Let S'_g be the surface obtained from S_g by removing an open disk. We consider the *relative mapping class group* $\text{Mod}(S'_g, \partial S'_g)$, which is $\pi_0(\text{Homeo}^+(S'_g, \partial S'_g))$, where $\text{Homeo}^+(S'_g, \partial S'_g)$ is the subgroup of elements of $\text{Homeo}^+(S'_g)$ that fix $\partial S'_g$ pointwise. Let $UT(S_g)$ denote the unit tangent bundle for S_g .

We have the following generalization of the classical Birman exact sequence due to Johnson [29].

Theorem 9.2 (Relative Birman exact sequence). *Let $g \geq 2$. We have*

$$1 \rightarrow \pi_1(UT(S_g)) \rightarrow \text{Mod}(S'_g, \partial S'_g) \rightarrow \text{Mod}(S_g) \rightarrow 1.$$

It was noticed by Johnson [29] that the entire image of $\pi_1(UT(S))$ lies in the Torelli subgroup of $\text{Mod}(S'_g, \partial S'_g)$. Therefore, the preimage of Γ_g in $\text{Mod}(S'_g, \partial S'_g)$ is a subgroup of the Torelli group $\mathcal{I}(S'_g, \partial S'_g)$. We call this group Γ'_{g+1} . By Theorem 9.1, Γ'_{g+1} is a Poincaré duality group of cohomological dimension $3(g+1) - 5$ (since $UT(S_g)$ is a closed aspherical 3-manifold, $\pi_1(UT(S_g))$ is a Poincaré duality group of dimension 3).

The inclusion of S'_g into S_{g+1} induces an injective homomorphism $\text{Mod}(S'_g, \partial S'_g) \rightarrow \text{Mod}(S_{g+1})$ (see, e.g., [38, Corollary 4.2]) which restricts to an inclusion on the level of Torelli groups. We define Γ_{g+1} to be the image of Γ'_{g+1} . Since Γ_{g+1} is isomorphic to Γ'_{g+1} , we have that Γ_{g+1} is a Poincaré duality group of dimension $3(g+1) - 5$, and this completes the proof of Theorem 1.1.

Remark. While the above argument certainly proves Theorem 1.1, we would like to point out that more is true. Not only are the Γ_g Poincaré duality groups, but they are also fundamental groups of closed aspherical manifolds, with base Γ_{g-1} and fiber $UT(S_{g-1})$. The proof of this stronger statement is given by Ivanov [26, Section 6.3].

Remark. Often the easiest way to get a lower bound on the cohomological dimension of a group is to find a large abelian subgroup (cf. Fact 6.5). It is a theorem of Vautaw [46], however, that the largest free abelian subgroup of $\mathcal{I}(S_g)$ has rank $2g - 3$. In Figure 9, we exhibit such a subgroup in genus 5 (the example generalizes to higher genus). In that example, each generator is a Dehn twist about a separating curve and so we see that we have \mathbb{Z}^{2g-3} inside $\mathcal{K}(S_g)$. This gives the lower bound for $\text{cd}(\mathcal{K}(S_g))$ given in the introduction.




FIGURE 9. The Dehn twists about these curves generate a free abelian subgroup of maximal rank in $\mathcal{I}(S_5)$.

Remark. The construction of the Mess subgroups clearly illustrates the discrepancy of g between $\text{cd}(\mathcal{I}(S_g))$ and $\text{vcd}(\text{Mod}(S_g))$. To get a group of cohomological dimension $4g - 5$ in $\text{Mod}(S_g)$, Mess simply augments the inductive construction of the Γ_g by adding the Dehn twist about a nonseparating curve in the new handle

at each stage. This new Dehn twist generates a direct factor of the Mess subgroup for $\text{Mod}(S_g)$. By only adding these Dehn twists on $g - k$ of the new handles, the construction gives the correct lower bound for Conjecture 1.2.




FIGURE 10

On the other hand, we note that the conjecture cannot be proven by applying the Quillen condition to the action of $\mathcal{I}^k(S_g)$ on $\mathcal{B}(S_g)$. To see this, we consider the group $\mathcal{I}^1(S_3)$, and the configuration shown in Figure 10; the handle being fixed homologically is the one spanned by a and b (say $x = [a]$). There is a 2-cell of $\mathcal{B}(S_3)$ corresponding to the c_i : $[c_1] + [c_2] + [c_3] = x$, and each of c_4 and c_5 adds a “borrowing dimension”. The conjecture says that $\text{cd}(\mathcal{I}^k(S_g)) = 6$, so we would need that the stabilizer of $\{c_i\}$ in $\mathcal{I}^1(S_3)$ has cohomological dimension at most 4. On the contrary, we can find a free abelian group of rank 5 in this stabilizer: the group generated by $T_{c_2}, T_{c_3}, T_{c_4}, T_{c_5}$, and T_d .

APPENDIX A. MULTITWISTS IN THE JOHNSON FILTRATION

This appendix gives a classification of the multitwists in each term of the Johnson filtration. We recall the notation from the introduction:

$$\Gamma = \pi_1(S_g), \quad \Gamma_1 = \Gamma, \quad \Gamma_{i+1} = [\Gamma, \Gamma_i], \quad \text{and} \\ \mathcal{N}_k(S_g) = \ker(\text{Mod}(S_g) \rightarrow \text{Out}(\Gamma/\Gamma_{k+1})).$$

Again, we have $\mathcal{N}_1(S_g) = \mathcal{I}(S_g)$ and $\mathcal{N}_2(S_g) = \mathcal{K}(S_g)$.

For a multicurve M , we again denote by $G(M)$ the free abelian group of rank $|M|$ generated by the Dehn twists in the curves of M .

Theorem A.1 answers the question of which multitwists lie in which terms of the Johnson filtration. The starting point is a theorem of Vautaw [46, Theorem 3.1], stated as part (1) of Theorem A.1, which gives the answer for $\mathcal{I}(S_g)$. Parts (2) and (3) of Theorem A.1 are new.

Theorem A.1. *Let M be a multicurve in S_g .*

- (1) *The group $G(M) \cap \mathcal{N}_1(S_g) = G(M) \cap \mathcal{I}(S_g)$ is generated by bounding pair maps $T_c T_d^{-1}$, where c and d are curves of M , and by Dehn twists about separating curves of M .*
- (2) *The group $G(M) \cap \mathcal{N}_2(S_g) = G(M) \cap \mathcal{K}(S_g)$ is generated by the Dehn twists about the separating curves of M .*
- (3) *For $k \geq 3$, the group $G(M) \cap \mathcal{N}_k(S_g)$ is trivial.*

Of course, there is nothing to say for the case $k = 0$. Also, as $\mathcal{I}(S_g)$ is trivial for $g \in \{0, 1\}$, there is nothing to do for those cases.

Dehn twists about separating curves are elements of $\mathcal{K}(S_g)$ (by definition); combining this with Statement (1) of Theorem A.1, Statement (2) of Theorem A.1 has the following equivalent formulation.

Theorem A.2. *If M is a multicurve in S_g consisting entirely of bounding pairs, then $G(M) \cap \mathcal{K}(S_g)$ is trivial.*

Using Statement (2) of Theorem A.1 and the fact that $\mathcal{N}_j(S_g) \leq \mathcal{N}_k(S_g)$ for $j \geq k$, we can recast Statement (3) as follows.

Theorem A.3. *If M is a multicurve in S_g consisting entirely of separating curves, then $G(M) \cap \mathcal{N}_3(S_g)$ is trivial.*

Punctures and boundary. Theorem A.1 also holds for surfaces that have one puncture or boundary component. For $k \geq 2$, the argument is essentially the same as for the closed case, but with the simplification that Fact A.7 is not needed. The case of the Torelli group is derived in a straightforward way from Vautaw’s theorem.

A.1. The Johnson kernel. In this section we prove Theorem A.2. Let M be a multicurve in S_g consisting entirely of nonseparating curves, and let f be an arbitrary multitwist supported in M ; without loss of generality, we assume that M is the canonical reduction system [7] for f ; that is, we assume that each curve of M is “used” by f . To prove the theorem, we need to show that f gives rise to a nontrivial element of $\text{Out}(\Gamma/\Gamma_3)$; that is, if φ is any representative of f that preserves the basepoint, then we can find an element ξ of $\pi_1(S_g)$ so that $\varphi_*(\xi)\xi^{-1}$ is not an element of Γ_3 .

We will require the following fact, whose proof we postpone until Section A.3.

Fact A.4. Suppose $g \geq 2$. Let x be an element of $\pi_1(S_g)$ represented by a nonseparating simple closed curve, and let y be an element of $\pi_1(S_g)$ that represents a nontrivial element \bar{y} of $H_1(S_g, \mathbb{Z})$ distinct from that of x . The commutator $[x, y]$ is not an element of $\Gamma_3 = [\Gamma, [\Gamma, \Gamma]]$.

Let $\{c, d\}$ be an *innermost* bounding pair in M , that is, a bounding pair that bounds a subsurface S' of S_g containing no curves of M in its interior (cf. the proof of Lemma 6.14). Note that S' must have positive genus.

Choose a regular neighborhood N of M , and choose a (standard) representative φ of f supported in N . Next, choose the basepoint for $\pi_1(S_g)$ to be an arbitrary point on the boundary component γ of N that is parallel to c and is not contained in S' . We consider γ (with arbitrary orientation) as an element of $\pi_1(S_g)$. Let α be any simple element of $\pi_1(S_g)$ that is homologically nontrivial, is homologically distinct from γ , and intersects N in a single arc (this arc contains the basepoint); it is possible to find such an element because S' has positive genus.

Since M is the canonical reduction system for f , there is a $k \neq 0$ so that

$$\varphi_*(\alpha) = \gamma^k \alpha \gamma^{-k}.$$

From here, one could check that $\varphi_*(\alpha)\alpha^{-1} = [\gamma^k, \alpha]$ is not an element of Γ_3 , and so φ does not give the trivial element of $\text{Aut}(\Gamma/\Gamma_3)$. However, this argument does not suffice to show that f gives a nontrivial element of $\text{Out}(\Gamma/\Gamma_3)$. Rather, we need to show that if we change φ_* by any inner automorphism, the result is still nontrivial in $\text{Aut}(\Gamma/\Gamma_3)$.

To this end, let β be a simple element of $\pi_1(S_g)$ that is homologically nontrivial and intersects N only at the basepoint; in this situation one can take $\beta = \gamma$ (in Section A.2, β will necessarily not equal γ). By construction φ_* fixes β .

Now, let δ be an arbitrary element of $\pi_1(S_g)$, let i_δ be the corresponding inner automorphism of $\pi_1(S_g)$, and let $\Psi = i_\delta \circ \varphi_*$ be the composite automorphism. We

have

$$\begin{aligned}\Psi(\alpha)\alpha^{-1} &= [\delta\gamma^k, \alpha], \text{ and} \\ \Psi(\beta)\beta^{-1} &= [\delta, \beta].\end{aligned}$$

The goal is to show that, for any choice of δ , at least one of these two commutators is not an element of $\Gamma_3 = [\Gamma, [\Gamma, \Gamma]]$. Since β represents a nontrivial element of $H_1(S_g, \mathbb{Z})$, it follows that $[\delta, \beta]$ is an element of Γ_3 if and only if δ is an element of Γ_2 , the commutator subgroup of $\pi_1(S_g)$; one direction is immediate from the definitions, and the other direction is Fact A.4. Thus, we may assume that $\delta \in \Gamma_2$, and we must show that $[\delta\gamma^k, \alpha]$ is not an element of Γ_3 . We use the following equality, which Magnus–Karrass–Solitar call a “Witt–Hall identity” [33, Theorem 5.1(10)]:

$$[\delta\gamma^k, \alpha] = \delta[\gamma^k, \alpha]\delta^{-1}[\delta, \alpha].$$

Since $\delta \in \Gamma_2$, the commutator $[\delta, \alpha]$ is, by definition, an element of Γ_3 . Also, by Fact A.4, and the fact that Γ_3 is normal in $\pi_1(S_g)$, the element $\delta[\gamma^k, \alpha]\delta^{-1}$ is not an element of Γ_3 . Thus, the commutator $[\delta\gamma^k, \alpha]$ is not an element of Γ_3 , and the proof is complete.

A.2. Higher terms. We now give the proof of Theorem A.3, which is completely analogous to the proof of Theorem A.2. Let M be a multicurve in S_g consisting entirely of separating curves, and let f be a multitwist supported in M ; again, we assume that M is the canonical reduction system for f . Let φ be a representative for f supported in a regular neighborhood N of M .

Let c be an innermost curve of M , by which we mean that there is a subsurface S' of S_g that is bounded by c and that contains no curves of M in its interior. Choose the basepoint for $\pi_1(S_g)$ to be an arbitrary point of the component of ∂N that is parallel to c but not contained in S' . Let γ be an element of $\pi_1(S_g)$ that travels around this component of ∂N once, and let α be an element of $\pi_1(S_g)$ that is contained in $S' \cup N$, and is represented by a nonseparating curve.

Take β to be a simple element of $\pi_1(S_g)$ that represents a nontrivial homology class and that intersects N only at the basepoint. Let δ be an arbitrary element of $\pi_1(S_g)$, let i_δ be the resulting inner automorphism, and let $\Psi = i_\delta \circ \varphi_*$. We have

$$\begin{aligned}\Psi(\alpha)\alpha^{-1} &= [\delta\gamma^k, \alpha], \text{ and} \\ \Psi(\beta)\beta^{-1} &= [\delta, \beta].\end{aligned}$$

As in the proof of Theorem A.2, at most one of these commutators is an element of Γ_4 , depending on whether or not δ is an element of Γ_3 . To make the analogous argument, we use the following fact, proven in Section A.3.

Fact A.5. Let x be an element of $\pi_1(S_g)$ represented by a nonseparating simple closed curve. If y is an element of $\pi_1(S_g)$ that is not homologous to x and is not an element of Γ_3 , then the commutator $[x, y]$ is not an element of Γ_4 .

One must also use the fact that simple elements of $\pi_1(S_g)$ (and their powers) are not elements of Γ_3 . This follows easily from Fact A.7 below.

A.3. The lower central series of a surface group. For an arbitrary group G , we denote by $\{G_i\}$ its lower central series, that is, $G_1 = G$, and $G_{i+1} = [G, G_i]$. Each group G_i/G_{i+1} is abelian, and we consider the graded abelian group

$$\bar{G} = G_1/G_2 \oplus G_2/G_3 \oplus G_3/G_4 \oplus \cdots .$$

The group \bar{G} has the structure of a graded Lie algebra over the ring \mathbb{Z} ; the Lie bracket is induced by the commutator operation in G . We denote by \bar{G}_i the graded piece of \bar{G} corresponding to G_i/G_{i+1} , and we will denote the group operation in \bar{G}_i by addition.

In this section, F denotes the free group on n letters x_1, \dots, x_n . We denote by \bar{x}_i the image of the generator x_i in \bar{F}_1 .

Fact A.6. The graded Lie algebra \bar{F} is the free Lie algebra over \mathbb{Z} with basis $\{\bar{x}_1, \dots, \bar{x}_n\}$.

Let $n = 2g$. The natural map $F \rightarrow \Gamma$ induces a map $\bar{F} \rightarrow \bar{\Gamma}$. Let $\omega \in \bar{F}_2$ be the element

$$\omega = [\bar{x}_1, \bar{x}_2] + \dots + [\bar{x}_{2g-1}, \bar{x}_{2g}].$$

We have the following fact, which is a special case of a theorem of Labute [32].

Fact A.7. The kernel of the natural map $\bar{F} \rightarrow \bar{\Gamma}$ is the ideal generated by ω .

Combining Facts A.6 and A.7, we see that $\bar{\Gamma}$ is isomorphic to the quotient of the free Lie algebra on $\bar{x}_1, \dots, \bar{x}_{2g}$ by the ideal generated by ω .

To prove Fact A.4, it is convenient to take x to be x_1 . By Fact A.7, we need to show that, in \bar{F} , the bracket $[\bar{x}_1, \bar{y}]$ is not equal to a power of ω . But if we write \bar{y} in terms of the \bar{x}_i -basis, and use the Witt–Hall formula

$$[\alpha + \beta, \gamma] = [\alpha, \gamma] + [\beta, \gamma],$$

it is evident that $[\bar{x}_1, \bar{y}]$ cannot equal a power of ω , since $g \geq 2$.

For Fact A.5, we can take $x = x_1$. If y is homologically nontrivial, then Fact A.4 applies. Therefore, we may assume that y corresponds to the nontrivial element $\bar{y} = \sum k_i [\bar{x}_i, \bar{x}_j]$ of \bar{F}_2 . Since $y \notin \Gamma_3$, we know that \bar{y} is not a power of ω . Again using the Witt–Hall identity, we see that the bracket $[\bar{x}_1, \bar{y}]$ is a nontrivial element of $\bar{\Gamma}_3$, and so $[x, y] \notin \Gamma_4$.

ACKNOWLEDGEMENTS

We would like to thank Ken Bromberg for many discussions on Teichmüller theory. We are also grateful to Tara Brendle, Benson Farb, Allen Hatcher, Chris Leininger, Justin Malestein, Andy Putman, Takuya Sakasai, Karen Vogtmann, and the referee for helpful comments and conversations. We would especially like to thank Tom Church for extensive discussions related to Lemma 6.16. The third author would like to thank the University of Utah for its support during this project.

REFERENCES

- [1] Problems in low-dimensional topology. In Rob Kirby, editor, *Geometric topology (Athens, GA, 1993)*, volume 2 of *AMS/IP Stud. Adv. Math.*, pages 35–473. Amer. Math. Soc., Providence, RI, 1997. MR1470751
- [2] Toshiyuki Akita. Homological infiniteness of Torelli groups. *Topology*, 40(2):213–221, 2001. MR1808217 (2001m:57022)
- [3] Mladen Bestvina, Kai-Uwe Bux, and Dan Margalit. Dimension of the Torelli group for $\text{Out}(F_n)$. *Invent. Math.*, 170(1):1–32, 2007. MR2336078 (2008k:57029)
- [4] Robert Bieri. Groupes à dualité de Poincaré. *C. R. Acad. Sci. Paris Sér. A-B*, 273:A6–A8, 1971. MR0284509 (44:1734)
- [5] Joan S. Birman. On Siegel’s modular group. *Math. Ann.*, 191:59–68, 1971. MR0280606 (43:6325)
- [6] Joan S. Birman. *Braids, links, and mapping class groups*. Princeton University Press, Princeton, N.J., 1974. Annals of Mathematics Studies, No. 82. MR0375281 (51:11477)

- [7] Joan S. Birman, Alex Lubotzky, and John McCarthy. Abelian and solvable subgroups of the mapping class groups. *Duke Math. J.*, 50(4):1107–1120, 1983. MR726319 (85k:20126)
- [8] A. Borel and J.-P. Serre. Corners and arithmetic groups. *Comment. Math. Helv.*, 48:436–491, 1973. Avec un appendice: Arrondissement des variétés à coins, par A. Douady et L. Hérault. MR0387495 (52:8337)
- [9] Kenneth S. Brown. *Cohomology of groups*, volume 87 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1982. MR672956 (83k:20002)
- [10] Peter Buser. *Geometry and spectra of compact Riemann surfaces*, volume 106 of *Progress in Mathematics*. Birkhäuser Boston Inc., Boston, MA, 1992. MR1183224 (93g:58149)
- [11] Thomas Church. Orbits of curves under the Johnson kernel. In preparation.
- [12] Thomas Church. Personal communication.
- [13] Max Dehn. *Papers on group theory and topology*. Springer-Verlag, New York, 1987. Translated from the German and with introductions and an appendix by John Stillwell, With an appendix by Otto Schreier. MR881797 (88d:01041)
- [14] Samuel Eilenberg and Tudor Ganea. On the Lusternik-Schnirelmann category of abstract groups. *Ann. of Math. (2)*, 65:517–518, 1957. MR0085510 (19:52d)
- [15] Benson Farb. Some problems on mapping class groups and moduli space. In *Problems on mapping class groups and related topics*, volume 74 of *Proc. Sympos. Pure Math.*, pages 11–55. Amer. Math. Soc., Providence, RI, 2006. MR2264130 (2007h:57018)
- [16] A. T. Fomenko, D. B. Fuchs, and V. L. Gutenmacher. *Homotopic topology*. Akadémiai Kiadó (Publishing House of the Hungarian Academy of Sciences), Budapest, 1986. Translated from the Russian by K. Mályusz. MR873943 (88f:55001)
- [17] Richard Hain. The rational cohomology ring of the moduli space of abelian 3-folds. *Math. Res. Lett.*, 9(4):473–491, 2002. MR1928867 (2003h:14068)
- [18] John L. Harer. Stability of the homology of the mapping class groups of orientable surfaces. *Ann. of Math. (2)*, 121(2):215–249, 1985. MR786348 (87f:57009)
- [19] John L. Harer. The virtual cohomological dimension of the mapping class group of an orientable surface. *Invent. Math.*, 84(1):157–176, 1986. MR830043 (87c:32030)
- [20] Allen Hatcher. Spectral sequences. Preliminary version, available at <http://www.math.cornell.edu/~hatcher/SSAT/SSATpage.html>.
- [21] Allen Hatcher. The cyclic cycle complex of a surface. arXiv:0806.0326.
- [22] Allen Hatcher. On triangulations of surfaces. *Topology Appl.*, 40(2):189–194, 1991. MR1123262 (92f:57020)
- [23] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002. MR1867354 (2002k:55001)
- [24] Y. Imayoshi and M. Taniguchi. *An introduction to Teichmüller spaces*. Springer-Verlag, Tokyo, 1992. Translated and revised from the Japanese by the authors. MR1215481 (94b:32031)
- [25] N. V. Ivanov. Complexes of curves and Teichmüller spaces. *Mat. Zametki*, 49(5):54–61, 158, 1991. MR1137173 (93a:32032)
- [26] Nikolai V. Ivanov. Mapping class groups. In *Handbook of geometric topology*, pages 523–633. North-Holland, Amsterdam, 2002. MR1886678 (2003h:57022)
- [27] Dennis Johnson. An abelian quotient of the mapping class group \mathcal{I}_g . *Math. Ann.*, 249(3):225–242, 1980. MR579103 (82a:57008)
- [28] Dennis Johnson. Conjugacy relations in subgroups of the mapping class group and a group-theoretic description of the Rochlin invariant. *Math. Ann.*, 249(3):243–263, 1980. MR579104 (82a:57007)
- [29] Dennis Johnson. The structure of the Torelli group. I. A finite set of generators for \mathcal{I} . *Ann. of Math. (2)*, 118(3):423–442, 1983. MR727699 (85a:57005)
- [30] Dennis Johnson. The structure of the Torelli group. II. A characterization of the group generated by twists on bounding curves. *Topology*, 24(2):113–126, 1985. MR793178 (86i:57011)
- [31] F. E. A. Johnson and C. T. C. Wall. On groups satisfying Poincaré duality. *Ann. of Math. (2)*, 96:592–598, 1972. MR0311796 (47:358)
- [32] John P. Labute. On the descending central series of groups with a single defining relation. *J. Algebra*, 14:16–23, 1970. MR0251111 (40:4342)
- [33] Wilhelm Magnus, Abraham Karrass, and Donald Solitar. *Combinatorial group theory*. Dover Publications Inc., Mineola, NY, second edition, 2004. Presentations of groups in terms of generators and relations. MR2109550 (2005h:20052)

- [34] Darryl McCullough and Andy Miller. The genus 2 Torelli group is not finitely generated. *Topology Appl.*, 22(1):43–49, 1986. MR831180 (87h:57015)
- [35] Greg McShane and Igor Rivin. A norm on homology of surfaces and counting simple geodesics. *Internat. Math. Res. Notices*, (2):61–69 (electronic), 1995. MR1317643 (96b:57014)
- [36] Geoffrey Mess. Unit tangent bundle subgroups of the mapping class groups. *Preprint IHES/M/90/30*, 1990.
- [37] Geoffrey Mess. The Torelli groups for genus 2 and 3 surfaces. *Topology*, 31(4):775–790, 1992. MR1191379 (93k:57003)
- [38] Luis Paris and Dale Rolfsen. Geometric subgroups of mapping class groups. *J. Reine Angew. Math.*, 521:47–83, 2000. MR1752295 (2001b:57035)
- [39] Andrew Putman. The Johnson homomorphism and its kernel. Preprint. arXiv:0904.0467.
- [40] Andrew Putman. Cutting and pasting in the Torelli groups. *Geom. Topol.*, 11:829–865 (electronic), 2007. MR2302503 (2008c:57049)
- [41] Jean-Pierre Serre. Cohomologie des groupes discrets. In *Prospects in mathematics (Proc. Sympos., Princeton Univ., Princeton, N.J., 1970)*, pages 77–169. Ann. of Math. Studies, No. 70. Princeton Univ. Press, Princeton, N.J., 1971. MR0385006 (52:5876)
- [42] Jean-Pierre Serre. *Trees*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003. Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation. MR607504 (82c:20083)
- [43] John R. Stallings. On torsion-free groups with infinitely many ends. *Ann. of Math. (2)*, 88:312–334, 1968. MR0228573 (37:4153)
- [44] Richard G. Swan. Groups of cohomological dimension one. *J. Algebra*, 12:585–610, 1969. MR0240177 (39:1531)
- [45] William P. Thurston. A norm for the homology of 3-manifolds. *Mem. Amer. Math. Soc.*, 59(339):i–vi and 99–130, 1986. MR823443 (88h:57014)
- [46] William R. Vautaw. Abelian subgroups of the Torelli group. *Algebr. Geom. Topol.*, 2:157–170 (electronic), 2002. MR1917048 (2003g:57027)

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF UTAH, 155 S 1400 EAST, SALT LAKE CITY, UTAH 84112-0090

E-mail address: `bestvina@math.utah.edu`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF VIRGINIA, KERCHOF HALL 229, CHARLOTTESVILLE, VIRGINIA 22903-4137

E-mail address: `kb2ue@virginia.edu`

DEPARTMENT OF MATHEMATICS, TUFTS UNIVERSITY, 503 BOSTON AVENUE, MEDFORD, MASSACHUSETTS 02155

E-mail address: `dan.margalit@tufts.edu`