


GENERATING THE TORELLI GROUP

ALLEN HATCHER AND DAN MARGALIT

ABSTRACT. We give a new proof of the theorem of Birman–Powell that the Torelli subgroup of the mapping class group of a closed orientable surface of genus at least 3 is generated by simple homeomorphisms known as bounding pair maps. The key ingredient is a proof that the subcomplex of the curve complex of the surface spanned by curves within a fixed homology class is connected.

1. INTRODUCTION

The *mapping class group* of a closed connected orientable surface S is $M(S) = \pi_0(\text{Diff}^+(S))$, the group of isotopy classes of orientation-preserving diffeomorphisms of S . Perhaps the simplest type of isotopically nontrivial diffeomorphism of S is a *Dehn twist* along an embedded closed curve. This is a diffeomorphism supported on an annular neighborhood of the curve, the effect of the diffeomorphism on arcs crossing the annulus being to twist these arcs around the annulus as shown in the following figure:



In the 1920s Dehn proved that $M(S)$ is generated by these twists, although this result was only published a decade later in [3, §10]; an English translation can be found in [4, Paper 8].

The group $M(S)$ has a natural action on $H_1(S) = H_1(S; \mathbb{Z})$, and the kernel of this action is known as the *Torelli group*, for which we use the notation $T(S)$. In this paper we will be interested in the analogue of Dehn's theorem for the Torelli group.

Certain Dehn twists belong to the Torelli group, namely the twists along curves that separate S . This can be seen by observing that for


2000 *Mathematics Subject Classification*. Primary: 20F36; Secondary: 57M07.

Key words and phrases. Torelli group, bounding pair maps.

The second author gratefully acknowledges support from the National Science Foundation and the Sloan Foundation.

each separating curve c , a basis for $H_1(S)$ can be chosen consisting of curves disjoint from c . Conversely, it is easy to see that if c is nonseparating, a twist along c acts nontrivially on the homology class of a curve that crosses c exactly once.

The next-simplest element of $T(S)$ after a separating twist is a *bounding pair map*, which is the composition of a twist along a nonseparating curve c and an inverse twist along another nonseparating curve d disjoint from c but representing the same homology class as c , so $c \cup d$ separates S into two subsurfaces having $c \cup d$ as their common boundary.



A bounding pair map, like a separating twist, can be realized by a motion of an embedding of S in \mathbb{R}^3 in which a subsurface is rotated through 360 degrees, a subsurface bounded by c in the case of a separating twist and by c and d in the case of a bounding pair map. One can see that a bounding pair map acts trivially on homology by noting that a basis for $H_1(S)$ can be chosen to consist of curves disjoint from c and d except for one curve that crosses each of c and d exactly once, and it is easy to see that this curve is taken to a homologous curve.

Theorem 1 (Birman–Powell). *The Torelli group $T(S)$ is generated by separating twists and bounding pair maps.*

Nontrivial separating twists exist only when the genus g of S is at least 2, and nontrivial bounding pair maps exist only when $g \geq 3$. Thus when $g = 1$ the Torelli group is trivial, a fact known long before the Birman–Powell theorem, and when $g = 2$ the theorem says that $T(S)$ is generated by separating twists. When $g \geq 3$, separating twists do not generate all of $T(S)$ but only a subgroup of infinite index known as the Johnson kernel [6]. On the other hand, it is easy to express separating twists as products of bounding pair maps when $g \geq 3$, as we recall in Proposition 11, so bounding pair maps alone generate $T(S)$ when $g \geq 3$. It is also easy to express all bounding pair maps in terms of those where the two curves that specify the map cut off a genus 1 subsurface of S ; see Proposition 12.

In this paper we give a proof of the Birman–Powell theorem that is in line with the standard proofs of Dehn’s theorem on $M(S)$ (see, e.g., [5, Theorem 4.1]). Dehn in fact found a finite set of twist maps that

generate $M(S)$, and the standard proofs of Dehn's theorem usually prove this as well. Finding a finite set of bounding pair maps that generate $T(S)$ when $g \geq 3$ is more difficult, and was first done by Johnson [8]. We do not attempt to reprove the finite generation here. Recently Putman [17] has greatly improved Johnson's finite generation result by showing that the number of generators can be reduced from Johnson's exponential function of g to a cubic function of g .

The genus 2 case of the Birman–Powell theorem was not stated explicitly in their original papers, but can easily be deduced from their methods. The genus 2 case is exceptional not only in the types of generators that are needed, but also in the fact that $T(S)$ is not finitely generated in this case, a result due originally to McCullough and Miller [12] and subsequently improved by Mess [13] who showed that $T(S)$ is a free group on a countably infinite set of twist generators. Another more recent proof of this can be found in [1] and we give a version of this proof at the end of Section 5.

Before sketching the idea for the new proof of the Birman–Powell theorem, let us say a few words about the two prior proofs in the literature.

First proof: Birman and Powell 1970s. This starts with a fact from group theory. Suppose we have a short exact sequence of groups

$$1 \longrightarrow A \longrightarrow B \xrightarrow{\pi} C \longrightarrow 1.$$

Let $\{b_1, \dots, b_n\}$ be a finite set of generators for B , and suppose we have a presentation for C in terms of the generating set $\{\pi(b_i)\}$. The relators for the presentation of C are words in the $\pi(b_i)$, and to each such relator w , the corresponding product of the b_i gives an element \tilde{w} of A if we identify A with the kernel of π . It is then a basic fact that the collection of all such \tilde{w} , together with their conjugates in B , forms a generating set for A (see the proof of [11, Theorem 2.1]).

In 1961 Klingen [10] gave an algorithm for finding a presentation of $\mathrm{Sp}(2g, \mathbb{Z})$ and ten years later Birman [2, Theorem 1] used this algorithm to give an explicit finite presentation for $\mathrm{Sp}(2g, \mathbb{Z})$. Birman's presentation therefore gave a generating set for $T(S)$ as above, but it was not immediately clear how to interpret the generators geometrically. In 1978, Powell recognized Birman's generators as Dehn twists about separating curves and bounding pair maps, or products of these [14, Theorem 2]. Neither the Birman paper nor the Powell paper contain complete details, as in both cases the required calculations are lengthy and technical.

Second proof: Putman 2007. Here the starting point is the fact that if a group G acts on a simply-connected simplicial complex X in such a way that the quotient of X by G is simply connected, then G is generated by stabilizers of simplices.

Putman [15] applied this strategy to the action of $T(S)$ on the curve complex $\mathcal{C}(S)$. This is the simplicial complex whose vertices are the isotopy classes of nontrivial simple closed curves in S , with simplices corresponding to collections of disjoint curves. Harer had shown that $\mathcal{C}(S)$ is $(2g - 3)$ -connected, hence simply connected when $g \geq 2$. Putman then showed that $\mathcal{C}(S)/T(S)$ is simply connected when $g \geq 2$.

By the fact about group actions stated above, this reduces the problem to understanding stabilizers of simplices of $\mathcal{C}(S)$. Just as in the proof that $M(S)$ is generated by Dehn twists, this step can be accomplished by induction on g using the Birman exact sequence.

A new proof. This also proceeds by induction on genus, but is based on a different fact about group actions on complexes that only requires the complexes to be connected. The complex we use is the subcomplex $\mathcal{C}_x(S)$ of $\mathcal{C}(S)$ spanned by curves that can be oriented so as to represent some fixed primitive class $x \in H_1(S)$. The key fact is therefore:

Theorem 2 (Putman). *For $g \geq 3$, the complex $\mathcal{C}_x(S)$ is connected.*

Putman's proof of this in [16, Theorem 1.9] uses Johnson's explicit finite generating set for $T(S)$, which in turn depends on the Birman–Powell theorem. The idea of the new approach is to reverse these dependencies to give a proof of Theorem 2 from scratch and deduce the Birman–Powell theorem from this.

The induction for the new proof of the Birman–Powell theorem starts with the case $g = 2$. The complex $\mathcal{C}_x(S)$ has dimension 0 in this case and is not connected, so instead we use a larger complex $\mathcal{B}_x(S)$ that appears in the proof of Theorem 2. It was shown in [1] that $\mathcal{B}_x(S)$ is contractible for all g , and we reprove this here. For $g \geq 3$ we only need that $\mathcal{B}_x(S)$ is connected, but when $g = 2$ the complex $\mathcal{B}_x(S)$ is 1-dimensional and we need that it is a tree so that we can use basic facts about groups acting on trees.

Our goal is to give a self-contained proof of the Birman–Powell theorem, and so this paper contains a number of proofs of known results, even when our proofs are not essentially different from the existing proofs.

Outline of the paper. In Section 2 we recall from [1] the construction of a complex $\mathcal{B}_x(S) \supset \mathcal{C}_x(S)$ whose points are isotopy classes

of oriented, weighted multicurves (collections of finitely many disjoint curves) in S representing the homology class x , and we show that $\mathcal{B}_x(S)$ is contractible. In Section 3 we prove Theorem 2 by showing the map $\pi_0\mathcal{C}_x(S) \rightarrow \pi_0\mathcal{B}_x(S)$ is injective. This is where most of the novelty of the paper occurs. The inductive step in the proof of the Birman–Powell theorem is given in Section 4. Finally, in Section 5 we complete the proof, using $\mathcal{B}_x(S)$ directly to handle the base case of the induction, genus 2.

Acknowledgments. We would like to thank Tara Brendle, Leah Childers, Thomas Church, John Etnyre, Chris Leininger, Andy Putman, and Saul Schleimer for helpful discussions.

2. REPRESENTING HOMOLOGY CLASSES BY MULTICURVES

In this section we reformulate some constructions from [1] involving the representations of elements of $H_1(S)$ by linear combinations of disjoint oriented curves in S .

By a *multicurve* in S we mean a collection, possibly empty, of finitely many disjoint simple closed curves in S , none of which bounds a disk in S and no two of which bound an annulus. Usually we will not distinguish between a multicurve and its isotopy class. If orientations are specified for each curve c_i in a multicurve $c = c_1 \cup \cdots \cup c_n$, then a linear combination $\sum_i k_i c_i$ with coefficients $k_i \in \mathbb{Z}$ determines a class $[\sum_i k_i c_i]$ in $H_1(S)$. If we allow coefficients $k_i \in \mathbb{R}$, then $\sum_i k_i c_i$ gives a class in $H_1(S; \mathbb{R})$. By reorienting the curves c_i if necessary we can assume $k_i \geq 0$ for each i . The linear combinations $\sum_i k_i c_i$ then correspond to points in the first orthant $[0, \infty)^n$ in \mathbb{R}^n . For each oriented multicurve c we have a corresponding orthant $O(c)$, and we can form a space $\mathcal{A}(S)$ by starting with the disjoint union of all such orthants $O(c)$, one for each isotopy class of oriented multicurves c , and then identifying the faces obtained by setting some coefficients k_i equal to 0 with the orthants corresponding to the multicurves obtained by deleting the corresponding curves c_i . (When c is empty, the orthant $O(c)$ reduces to just the origin in \mathbb{R}^0 , so the origins of all the orthants $O(c)$ are identified.)

The natural map $h : \mathcal{A}(S) \rightarrow H_1(S; \mathbb{R})$ sending a weighted oriented multicurve $\sum_i k_i c_i$ to its homology class is linear on each orthant $O(c)$. For a nonzero class $x \in H_1(S; \mathbb{R})$ we define $\mathcal{A}_x(S) = h^{-1}(x)$. This is a cell complex whose cells are the intersections of orthants $O(c) = [0, \infty)^n$ with affine planes in \mathbb{R}^n . These ‘cells’ $E(c) = O(c) \cap h^{-1}(x)$ can be noncompact, so $\mathcal{A}_x(S)$ will not be a cell complex in the usual sense, but something more general. To guarantee that the cells $E(c)$ are compact

we need to impose a further condition on the oriented multicurves c , namely that if we translate the affine plane that determines $E(c)$ until this plane passes through the origin, which amounts to taking $O(c) \cap h^{-1}(0)$, then $O(c) \cap h^{-1}(0) = \{0\}$. In other words, no nontrivial linear combination $\sum_i k_i c_i$ with each $k_i \geq 0$ represents 0 in $H_1(S; \mathbb{R})$. Taking only orthants $O(c)$ for c satisfying this extra condition yields a subspace $\mathcal{B}(S)$ of $\mathcal{A}(S)$. The corresponding subspace $\mathcal{B}_x(S)$ of $\mathcal{A}_x(S)$ has the structure of a cell complex in the usual sense, with cells $E(c)$ that are compact convex polyhedra.

There is another way to characterize the compact cells $E(c)$ that is somewhat more geometric:

Proposition 3. *A cell $E(c)$ is compact if and only if no submulticurve of c with the induced orientation from c is a boundary, representing 0 in $H_1(S)$.*

We call an oriented multicurve satisfying this property *reduced*. As we will see in the proof, an equivalent condition is that no submulticurve with its induced orientation is the oriented boundary of an oriented subsurface of S .

Proof. As noted earlier, $E(c)$ is compact if and only if no nontrivial linear combination $\sum_i k_i c_i$ with each $k_i \geq 0$ is trivial in $H_1(S)$. This obviously implies that no submulticurve is a boundary, either in the homological sense or the geometric sense of bounding an oriented subsurface of S .

For the converse, suppose that some nontrivial sum $\sum_i k_i c_i$, with each $k_i \in [0, \infty)$, is a boundary, say $\sum_i k_i c_i = \partial(\sum_j l_j R_j)$ where the R_j 's are the closures of the components of $S - c$, oriented via a fixed orientation of S . Since $\partial(\sum_j R_j) = 0$, we can add a large constant l to each l_j to guarantee that $l_j > 0$ for all j . Let R be the union of the R_j 's with maximal l_j . Since $\sum_i k_i c_i$ is nontrivial, the surface R is a proper subsurface of S . Then the equation $\sum_i k_i c_i = \partial(\sum_j l_j R_j)$ implies that ∂R is a nonzero linear combination of the oriented curves c_i with each coefficient equal to 0 or 1. Thus c has a null-homologous submulticurve, in fact a submulticurve that bounds a subsurface of S . \square

A multicurve c has a dual graph $G(c)$ whose vertices correspond to components of $S - c$ and whose edges correspond to components of c . If c is an oriented multicurve, then $G(c)$ becomes an oriented graph by fixing an orientation of S and a rule for passing from an orientation of c to a transverse orientation. The condition for c to be reduced can be translated into a condition on $G(c)$:

Proposition 4. *An oriented multicurve c is reduced if and only if its dual graph $G(c)$ is recurrent: Every edge of $G(c)$ lies in a loop consisting of a finite sequence of edges traversed in the directions given by their orientations in $G(c)$. Such loops can be assumed to be embedded.*

This recurrence condition can be restated in terms of c as saying that through every point of c there passes a closed oriented loop in S transverse to c whose algebraic and geometric intersection numbers with c are equal, so the loop always crosses c in the same direction.

Proof. If c contains a bounding submulticurve c' , then any closed loop in S intersecting c' transversely must cross c' in both directions, so $G(c)$ is not recurrent. Conversely, suppose $G(c)$ is not recurrent. Let \overline{G} be the quotient graph of $G(c)$ obtained by collapsing to a point each component of the subgraph consisting of edges that lie in closed loops of oriented edges. Then \overline{G} contains no such edges. Hence \overline{G} must contain at least one “sink” vertex whose abutting edges are all oriented toward the vertex. These edges correspond to a bounding submulticurve of c . \square

Proposition 5. *The cell $E(c)$ of $\mathcal{B}_x(S)$ corresponding to an oriented multicurve c has dimension equal to one less than the number of connected components of $S - c$.*

Proof. Consider cellular homology with coefficients in \mathbb{R} for a cell structure on S containing c as a subcomplex. With notation as in the proof of Proposition 3, the regions R_j generate a subgroup $\langle R_j \rangle$ of the 2-chains and the curves c_i generate a subgroup $\langle c_i \rangle$ of the 1-chains. The boundary map $\partial : \langle R_j \rangle \rightarrow \langle c_i \rangle$ has 1-dimensional kernel $H_2(S; \mathbb{R})$ so its image has dimension one less than the number of R_j 's. Cosets of this image are the planes $h^{-1}(x)$, and intersecting one of these with the orthant $O(c)$ gives the cell $E(c)$, with the stated dimension. \square

Up until this point the class x was any nonzero element of $H_1(S; \mathbb{R})$, but from now on we restrict attention to classes in $H_1(S) = H_1(S; \mathbb{Z})$.

Proposition 6. *For a nonzero class $x \in H_1(S)$, the coefficients of a vertex $\sum_i k_i c_i$ of $\mathcal{B}_x(S)$ are integers.*

Proof. By Proposition 5 the multicurve $c = \cup_i c_i$ has connected complement, so for each c_j there is a transverse curve d_j intersecting c_j once and disjoint from the other c_i 's. The algebraic intersection number of d_j with $\sum_i k_i c_i$ is then $\pm k_j$, so if $\sum_i k_i c_i$ represents an integral homology class, k_j must be an integer. \square

Proposition 7. *Let $g \geq 1$, and let $x \in H_1(S)$ be any nonzero element. Then the complexes $\mathcal{A}_x(S)$ and $\mathcal{B}_x(S)$ are contractible.*

Proof. We follow the plan of the second proof in [1]. First consider $\mathcal{A}_x(S)$. To prove that $\mathcal{A}_x(S)$ is contractible we will construct a canonical “linear” path in $\mathcal{A}_x(S)$ joining any two given points α and β , assuming that the multicurves underlying α and β have first been isotoped to intersect transversely with the minimum number of points of intersection. With this minimality condition the configuration formed by the union of the two multicurves is unique up to isotopy of S , which will ensure that the construction is well defined on isotopy classes.

To a weighted multicurve $\alpha = \sum_i k_i c_i$ representing a point in $\mathcal{A}_x(S)$ we associate a map $f_\alpha : S \rightarrow S^1$ in the following way. First, choose disjoint product neighborhoods $c_i \times [0, k_i]$ of the curves c_i in S . (We can assume each $k_i > 0$ by deleting any c_i with $k_i = 0$.) From these product neighborhoods we obtain a quotient map q from S to the graph $G(c)$ by projecting each $c_i \times [0, k_i]$ to $[0, k_i]$ and then to the corresponding edge of $G(c)$, with the complementary components of the thickened c in S mapping to the corresponding vertices of $G(c)$. The weights k_i determine lengths for the edges of $G(c)$ making it into a metric graph, with edges oriented via the orientation of c . There is then a natural map $\varphi : G(c) \rightarrow S^1 = \mathbb{R}/\mathbb{Z}$ defined up to rotations of S^1 by the condition that it is an orientation-preserving local isometry on each edge of $G(c)$. Namely, choose a vertex of $G(c)$ and send it to an arbitrary point in \mathbb{R} . This determines a map on adjacent edges sending them isometrically to \mathbb{R} preserving orientations, then continue inductively for edges adjacent to the previous edges. Loops in $G(c)$ have signed length equal to the algebraic intersection number of lifted loops in S with x , and these intersection numbers are integers since x is an integral homology class, so when we pass to the quotient \mathbb{R}/\mathbb{Z} we have a well-defined map $\varphi : G(c) \rightarrow S^1$. Changing the initial vertex or its image in \mathbb{R} has the effect of composing φ with a rotation of S^1 . The composition φq is then a map $f_\alpha : S \rightarrow S^1$. This corresponds to the class in $H^1(S; \mathbb{Z})$ Poincaré dual to x . We can arrange that f_α is a smooth map by parametrizing the annuli $c_i \times [0, k_i]$ suitably.

For a second point β in $\mathcal{A}_x(S)$ we choose annular neighborhoods of its curves that intersect the neighborhood of α in rectangles around the points where α and β intersect, and then we construct the associated function f_β by the same procedure as for f_α . We would like to define a one-parameter family of functions $S \rightarrow S^1$ by the formula $(1-t)f_\alpha + tf_\beta$ for $0 \leq t \leq 1$. This does not quite make sense as it stands since scalar multiplication is not defined for maps to S^1 , but we can give

it meaning by considering the covering space \tilde{S} of S corresponding to the kernel of the map $\pi_1(S) \rightarrow \pi_1(S^1)$ induced by f_α and f_β , which are homotopic since α and β both represent x . Then f_α and f_β lift to maps \tilde{f}_α and \tilde{f}_β from \tilde{S} to \mathbb{R} , and we can form the linear combination $(1-t)\tilde{f}_\alpha + t\tilde{f}_\beta$. This is equivariant with respect to the action of \mathbb{Z} as deck transformations in \tilde{S} and \mathbb{R} , so it passes to a well-defined map $f_t = (1-t)f_\alpha + tf_\beta$ from S to S^1 .

The critical points of f_t are the closures of the components of the complement of the union of the annular neighborhoods of α and β . In the interiors of the rectangles where these neighborhoods intersect there are no critical points since the gradient vectors of f_t are the vectors $(1-t)\nabla f_\alpha + t\nabla f_\beta$ which are nonzero. Since there are finitely many complementary components of $\alpha \cup \beta$, the function f_t has finitely many critical values, each of which varies linearly with t . For fixed t the complement of the critical values consists of finitely many open intervals in S^1 , and the preimages of these intervals consist of finitely many open annuli in S , thickenings of disjoint curves which are oriented transversely by ∇f_t , with weights given by the lengths of the corresponding intervals in S^1 . These curves determine a weighted oriented multicurve representing a point α_t in $\mathcal{A}_x(S)$ by discarding any trivial curves and replacing isotopic curves by a single curve, weighted by the appropriate signed sum of the weights of the isotopic curves. For $t = 0$ we have $\alpha_0 = \alpha$ and for $t = 1$ we have $\alpha_1 = \beta$. The point $\alpha_t \in \mathcal{A}_x(S)$ varies continuously with t since the functions f_t vary smoothly and the intervals of noncritical values vary continuously, shrinking to length zero when critical values coalesce. This happens only finitely often for the path α_t since the finitely many critical values are varying linearly with t . By similar reasoning the path α_t varies continuously with the weights on the original multicurve α . Thus by fixing β and letting α vary over all of $\mathcal{A}_x(S)$ we obtain a contraction of $\mathcal{A}_x(S)$.

Now we show that $\mathcal{A}_x(S)$ deformation retracts onto $\mathcal{B}_x(S)$, which implies that $\mathcal{B}_x(S)$ is also contractible. The procedure here will be the same “draining” process as in [1]. If a point $\sum_i k_i c_i$ in $\mathcal{A}_x(S)$ is not reduced, let $\{R_j\}$ be the collection of oriented compact subsurfaces of S whose oriented boundary is a subset of the c_i 's, respecting their given orientations. We deform $\sum_i k_i c_i$ by subtracting $t \sum_j \partial R_j$ for increasing $t \geq 0$ until one or more c_i 's becomes 0. Deleting these c_i 's and the R_j 's whose boundaries include these c_i 's, we then iterate the process until we obtain a reduced weighted multicurve in $\mathcal{B}_x(S)$. It is clear this process depends continuously on the initial point $\sum_i k_i c_i$ and so defines a deformation retraction of $\mathcal{A}_x(S)$ into $\mathcal{B}_x(S)$. \square

Canonical triangulations of $\mathcal{A}_x(S)$ and $\mathcal{B}_x(S)$. Although we will not need this in the rest of the paper, there is a canonical subdivision of the cell complex $\mathcal{B}_x(S)$ as a simplicial complex whose vertices are the integer points of $\mathcal{B}_x(S)$, the linear combinations $\sum_i k_i c_i$ with positive integer coefficients k_i . This subdivision can be obtained as follows. In the preceding proof we associated to $\sum_i k_i c_i$ a map $f : S \rightarrow S^1$. This factors through the oriented metric graph $G(c)$ associated to $\sum_i k_i c_i$, with an induced map $\varphi : G(c) \rightarrow S^1$. Recall that the dimension of the cell $E(c)$ is 1 less than the number of complementary regions of c , which is the number of vertices of $G(c)$. This cell is subdivided by the various hyperplanes where two vertices of $G(c)$ have the same image under φ , hyperplanes defined by linear equations in the variables k_i with integer coefficients and integer constant terms. These hyperplanes subdivide $E(c)$ into simplices whose barycentric coordinates are the lengths of the segments of $S^1 = \mathbb{R}/\mathbb{Z}$ between adjacent images of vertices of $G(c)$. The vertices of the subdivision of $E(c)$ are thus the points $\sum_i k_i c_i$ where all vertices of $G(c)$ have the same image under φ , which is equivalent to saying that all the coefficients k_i are integers.

The same procedure works more generally for $\mathcal{A}_x(S)$, where the non-compact cells are subdivided into infinitely many simplices.

3. THE COMPLEX OF HOMOLOGOUS CURVES

In this section we prove Theorem 2, that $\mathcal{C}_x(S)$ is connected when $g \geq 3$ and x is any nonzero primitive class in $H_1(S)$. Recall the basic fact (see, e.g., [5, Proposition 6.2]) that primitive classes x are exactly those represented by oriented simple closed curves with coefficient 1. Thus when x is primitive, $\mathcal{C}_x(S)$ is the subcomplex of $\mathcal{B}_x(S)$ consisting of the cells that are simplices corresponding to cycles $\sum_i k_i c_i$ where the c_i 's are disjoint oriented curves each representing the homology class x and $\sum_i k_i = 1$. To prove that $\mathcal{C}_x(S)$ is connected it will suffice to show that the map $\pi_0 \mathcal{C}_x(S) \rightarrow \pi_0 \mathcal{B}_x(S)$ is injective when $g \geq 3$ since we already know that $\mathcal{B}_x(S)$ is connected. We will show in fact that each edge path γ in $\mathcal{B}_x(S)$ with endpoints in $\mathcal{C}_x(S)$ is homotopic, fixing endpoints, to an edge path in $\mathcal{C}_x(S)$.

Deforming the edge path γ into $\mathcal{C}_x(S)$ will be done by a sequence of local deformations, gradually decreasing the maximum value along γ of the “weight function” $W : \mathcal{B}_x(S) \rightarrow (0, \infty)$ defined by


$$W\left(\sum_i k_i c_i\right) = \sum_i k_i$$

In terms of dual graphs, W measures the total length of all the edges. The function W is linear on cells of $\mathcal{B}_x(S)$ and takes integer values on vertices. It follows that the image of W is contained in $[1, \infty)$. As x is

assumed to be primitive, we have $\mathcal{C}_x(S) = W^{-1}(1)$ since W takes the value 1 on vertices of $\mathcal{C}_x(S)$, hence on simplices of $\mathcal{C}_x(S)$, and if $W = 1$ on all vertices of a cell of $\mathcal{B}_x(S)$ then these vertices lie in $\mathcal{C}_x(S)$ and span a simplex of $\mathcal{C}_x(S)$.

Thus it will suffice to deform γ to decrease the maximum value of W along its vertices to 1. There will be two main steps. First, when the maximum occurs at two successive vertices of γ we will deform this edge of γ across a 2-cell in $\mathcal{B}_x(S)$ having smaller values of W on all the other vertices of the cell. Then by a more complicated procedure we will deform γ on the two edges surrounding a vertex where W is maximal to decrease the maximum along this part of γ .

Proof of Theorem 2. Let $\langle v_0, v_1 \rangle$ be an edge of $\mathcal{B}_x(S)$ joining vertices v_0 and v_1 . Associated to the points v_t along this edge, $0 \leq t \leq 1$, are dual graphs G_t . These are oriented metric graphs having two vertices a and b for $0 < t < 1$, so they have the form shown in the figure below. Since G_t is recurrent, there is at least one edge from a to b and at least one edge from b to a , but the number of edge-loops at a or b can be zero.



As t varies from 0 to 1, the lengths of the edges joining a and b vary, with all edges from a to b increasing in length at the same rate that all edges from b to a shrink, or vice versa. This corresponds to varying the weights in the weighted oriented multicurve v_t by adding a multiple of the boundary of the subsurface of S corresponding to a or b . If there are edge-loops at a or b , their lengths do not change. When t reaches 0 or 1, at least one edge joining a and b shrinks to length 0 and the vertices a and b coalesce.

Suppose that G_t has at least two edges entering at vertex a . We can pinch equal-length segments of two of these edges together at a to produce a new metric graph G'_t with three vertices. Assuming that the subsurface of S corresponding to a is not simply a pair of pants, we can realize G'_t as the dual graph for a point in a 2-cell of $\mathcal{B}_x(S)$ where the subsurface of S corresponding to the new vertex is a pair of pants, two of whose boundary curves correspond to the two edges of G_t being pinched together, with the third boundary curve corresponding to the new edge of G'_t . Note that pinching G_t to G'_t preserves the recurrence property so we do indeed have a 2-cell of $\mathcal{B}_x(S)$. In similar fashion we

could realize the graph G'_t obtained by pinching segments of two edges of G_t exiting a , or two edges entering or exiting b .

Let us apply this construction when $\langle v_0, v_1 \rangle$ is an edge of γ with $W(v_0) = W(v_1)$ and this is the maximum value of W along γ . The condition $W(v_0) = W(v_1)$ means that at both a and b there will be the same number of entering edges as exiting edges since varying t does not change the total length of all the edges of G_t . If the number of entering and exiting edges is equal to 1 at both a and b then G_t is just a circle and $W(v_0) = W(v_1) = 1$ since the class x is primitive. We may thus assume the number of entering and exiting edges at a , say, is greater than 1. We can then perform the pinching operation at a and this decreases the total length of G_t , so we obtain a 2-cell with the function W taking on its maximum value only along the edge $\langle v_0, v_1 \rangle$ of this 2-cell, since the 2-cell is a convex polygon and W is a nonconstant linear function on this polygon. If we modify γ by pushing the edge $\langle v_0, v_1 \rangle$ across this 2-cell to the complementary edges in the boundary of the cell, we have then improved the situation so that W has strictly smaller values between v_0 and v_1 . After repeating this step finitely many times we can arrange that the maximum value of W along γ occurs only at isolated vertices.

Now let $\langle v_0, v_1 \rangle$ be an edge of γ with $W(v_0)$ an isolated maximum of W along γ . A special case is when one of the vertices a or b of the graph G_t for points in the interior of this edge has valence 3 and the corresponding subsurface of S is a pair of pants. In this case we call $\langle v_0, v_1 \rangle$ a *P-edge*. We wish to reduce to the case that all edges $\langle v_0, v_1 \rangle$ of γ adjacent to vertices v_0 with maximal W value are P-edges, so suppose on the contrary that $\langle v_0, v_1 \rangle$ is not a P-edge. Since $W(v_1) < W(v_0)$, the number of edges in G_t from a to b will be different from the number of edges from b to a . For whichever type of edge there are more of, we can pinch two of the edges of this type together at either a or b , and this gives rise to a deformation of $\langle v_0, v_1 \rangle$ across a 2-cell of $\mathcal{B}_x(S)$ as in the preceding paragraph since we assume $\langle v_0, v_1 \rangle$ is not a P-edge. Then for the new path γ the new edge at v_0 is a P-edge. Iterating this step, we can arrange that all edges of γ at vertices with W maximum are P-edges.

Let v_0 be a vertex of γ with W maximal and greater than 1. If we cut S along the multicurve given by v_0 we obtain a cobordism R between two copies of this multicurve, which we label $\partial_+ R$ and $\partial_- R$. A P-edge from v_0 then corresponds to a pair of pants in R with two boundary curves in $\partial_+ R$ or two boundary curves in $\partial_- R$. If the genus of S is at least 3, such a pair of pants is uniquely determined by its third boundary curve, which gives a vertex in the curve complex $\mathcal{C}(R)$.

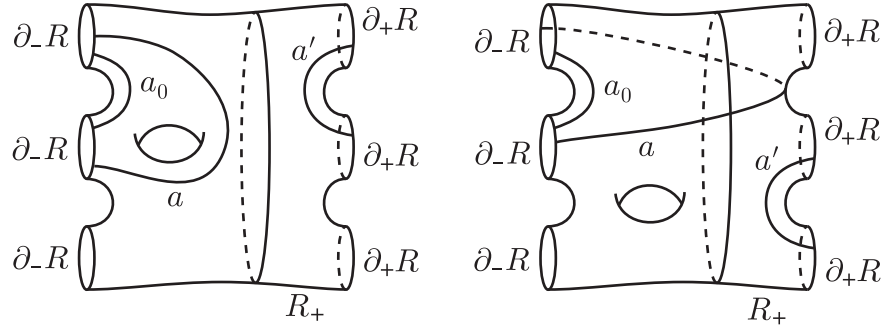
Such vertices span a subcomplex $\mathcal{P}(R)$ of $\mathcal{C}(R)$. We will prove below that $\mathcal{P}(R)$ is connected when the genus of S is at least 3. Assuming this, we finish the proof of the theorem as follows. By the preceding paragraph, we can assume the edges in γ on both sides of v_0 are P-edges. Since $\mathcal{P}(R)$ is connected, we can interpolate between these two P-edges a sequence of P-edges, each corresponding to a pair of pants disjoint from the next one. Each pair of successive P-edges then forms two adjacent edges of a 2-cell of $\mathcal{B}_x(S)$ (either a triangle or a square). We then deform γ by pushing across each of these 2-cells, thereby replacing the two edges of γ adjacent to v_0 by a sequence of edges along which W has values smaller than $W(v_0)$ since W is linear on each of the 2-cells.

After finitely many iterations of these steps we eventually deform γ , staying fixed on its endpoints, to a path in $W^{-1}(1) = \mathcal{C}_x(S)$. \square

Lemma 8. *The complex $\mathcal{P}(R)$ is connected when $g \geq 3$.*

Proof. Instead of regarding vertices of $\mathcal{P}(R)$ as isotopy classes of pairs of pants in R we can regard them as isotopy classes of arcs in R joining two curves of ∂_-R or two curves of ∂_+R , where the pair of pants corresponding to such an arc is a thickening of the union of the arc and the two curves at its endpoints. When we say ‘‘arc’’ in what follows, we will mean an arc giving a vertex of $\mathcal{P}(R)$ in this way. A simplex of $\mathcal{P}(R)$ corresponds to a collection of disjoint arcs joining disjoint pairs of curves of ∂R . Let us fix a standard arc a_0 joining two curves in ∂_-R . An arbitrary arc a can be connected to an arc joining the same two curves of ∂_-R as a_0 by a sequence of at most two edges of $\mathcal{P}(R)$, first by choosing an arc a' disjoint from a joining curves of ∂_+R if a does not already do this, then by choosing an arc disjoint from a' joining the two curves of ∂a_0 . Thus it suffices to connect an arbitrary arc a joining the two curves of ∂_-R containing ∂a_0 to the arc a_0 by a sequence of edges in $\mathcal{P}(R)$.


First we do this for three special types of arcs. For the first two cases we fix a genus 0 subsurface R_+ of $R - a_0$ containing ∂_+R and having just one more boundary curve, which lies in the interior of R .



(i) An arc a that is disjoint from R_+ as in the first figure above. In this case there is an edge joining a to an arc a' in R_+ and then an edge joining a' to a_0 .

(ii) An arc a that intersects R_+ in a single subsegment that separates R_+ into two components, one of which contains only one curve of $\partial_+ R$ and the other of which contains at least two curves of $\partial_+ R$, assuming that $\partial_+ R$ has at least three curves in total. This case is illustrated in the second figure above. In this case we can choose an arc a' in R_+ disjoint from a and proceed as in (i).

(iii) An arc a as in the figure below, in the case that $\partial_+ R$ contains just two curves, hence R has genus at least 1 since $g \geq 3$. Then there are edges joining a to a' and then to a_0 .



Now we reduce the general case to these special cases. Let us regard the curve of $\partial_- R$ at one end of a_0 as a puncture p rather than a boundary component. If we allow this puncture to move around anywhere in the surface R' obtained from R by filling in this puncture (or equivalently, collapsing the boundary component of R to a point), then any other arc a with the same endpoints as a_0 can be isotoped to a_0 . This implies that a_0 can be transformed to a (or an arc isotopic to a) by a diffeomorphism h of R obtained by dragging p around a loop in R based at p . Such diffeomorphisms form a subgroup of $M(R', p)$, the image of the boundary map $d : \pi_1(R', p) \rightarrow \pi_0 \text{Diff}^+(R', p)$ in the long exact sequence of homotopy groups associated to the fibration $\text{Diff}^+(R', p) \rightarrow \text{Diff}^+(R') \rightarrow R'$ obtained by evaluating diffeomorphisms

of R' at p , with fiber the subgroup $\text{Diff}^+(R', p)$ of $\text{Diff}^+(R')$ consisting of diffeomorphisms that fix p . The map d associates to a loop at p the diffeomorphism obtained by dragging p around this loop. (The long exact sequence in fact reduces to a short exact sequence, the Birman exact sequence.)

Since d is a homomorphism, it follows that the diffeomorphism h is isotopic to a composition $h_1 \cdots h_n$ of diffeomorphisms h_i obtained by dragging p around a suitable sequence of loops that generate $\pi_1(R', p)$. For such generators we can choose loops in $R - R_+$, producing arcs $a = h_i(a_0)$ as in (i), or loops that wind once around a single curve of $\partial_+ R$, producing arcs $a = h_i(a_0)$ as in (ii) or (iii). By the special cases (i-iii) there is an edge path in $\mathcal{P}(R)$ joining a_0 to $h_i(a_0)$, for each i . By applying the product $h_1 \cdots h_{i-1}$ to this edge path we obtain an edge path joining $h_1 \cdots h_{i-1}(a_0)$ to $h_1 \cdots h_i(a_0)$. Stringing these edge paths together, we obtain an edge path from a_0 to the arc $h(a_0) = a$. \square

4. THE INDUCTIVE STEP

We will prove Theorem 1 by induction on genus. In this section we give the inductive step, deferring the base case of genus 2 until the next section since it requires methods that are special to that case.

The inductive step will use the following basic fact about group actions:

Suppose a group G acts on a connected cell complex X . Let $A \subseteq G$ be a subset with the property that, for any two vertices of X connected by an edge, there is an element of A taking one vertex to the other. Then G is generated by the union of A and the set of vertex stabilizers.


Since $\mathcal{C}_x(S)$ is connected when $g \geq 3$, we may apply this fact to the case of the $T(S)$ action on $\mathcal{C}_x(S)$. For A we choose the set of bounding pair maps and twists about separating curves in $T(S)$. The condition on edges is verified in the next lemma. Here and in what follows we use the notation T_a for the twist along the curve a .

Lemma 9. *If v and w are vertices of $\mathcal{C}_x(S)$ that are connected by an edge, then there is a bounding pair map $T_a T_b^{-1}$ in $T(S)$ with $T_a T_b^{-1}(v) = w$.*

It is worth pointing out that this lemma together with the connectedness of $\mathcal{C}_x(S)$ immediately implies the non-obvious fact, known to Johnson [7, page 253, line 6], that when $g \geq 3$, any two oriented curves in S in the same homology class are equivalent under the action of the

Torelli group. This also holds when $g = 2$, as we will see in the next section.

Proof. We can view S as a torus with handles attached, with v and w as longitudes on the torus. Then if we choose a and b as meridians on the torus as shown in the figure below, it is clear that $T_a T_b^{-1}$ takes v to w . \square



To finish the inductive step it remains to show that for each vertex v of $\mathcal{C}_x(S)$, the stabilizer of v in $T(S)$ is contained in the subgroup of $T(S)$ generated by bounding pair maps and twists about separating curves. If v is represented by an oriented curve a in the homology class x , the stabilizer of v in $T(S)$ is the subgroup $T(S, a)$ represented by diffeomorphisms leaving a invariant. More generally, let $M(S, a)$ denote the stabilizer of a in $M(S)$. There is a natural homomorphism $\eta : M(S, a) \rightarrow M(S', P)$ where S' is the closed surface of genus $g - 1$ obtained from S by cutting along a and collapsing the resulting two boundary curves to a pair $P = \{p, q\}$ of distinguished points, and $M(S', P)$ is the mapping class group of S' fixing each of these two points. The kernel of η is the infinite cyclic group generated by T_a [5, Proposition 3.20]. Since $H_1(S', P) = H_1(S, a) = H_1(S)/\langle [a] \rangle$, and since $T_a^k \notin T(S)$ for $k \neq 0$, it follows that η restricts to an isomorphism:

$$\eta : T(S, a) \rightarrow T(S', P).$$

where the latter group is the kernel of the natural homomorphism $M(S', P) \rightarrow \text{Aut}(H_1(S', P))$.

If T_d is a twist about a separating curve d in $S' - P$, then $\eta^{-1}(T_d)$ is either a bounding pair map $T_d T_a^{-1}$ or a twist about a separating curve T_d , depending on whether or not d separates the two points of P in S' . If $T_c T_d^{-1}$ is a bounding pair map in $T(S', P)$, then $c \cup d$ does not separate the two points of P , otherwise $T_c T_d^{-1}$ would not preserve the homology class of an arc that intersects c once and is disjoint from d . It follows that $\eta^{-1}(T_c T_d^{-1})$ is a bounding pair map in $T(S, a)$. Thus in order to show that $T(S, a)$ is generated by twists about separating

curves and bounding pair maps it suffices to show that $T(S', P)$ has this property.

We will prove this by considering two short exact sequences:

$$\begin{aligned} 1 \rightarrow K_1 \rightarrow T(S', p) \rightarrow T(S') \rightarrow 1 \\ 1 \rightarrow K_2 \rightarrow T(S', P) \rightarrow T(S', p) \rightarrow 1 \end{aligned}$$

These are analogs for the Torelli group of Birman exact sequences for the full mapping class group; the first was considered by Johnson [8, Lemma 3] and the second by van den Berg [18, Proposition 2.4.1] and Putman [15, Theorem 4.1].

Consider first the first sequence. Here $T(S', p)$ is the subgroup of $M(S', p)$ acting trivially on $H_1(S', p) = H_1(S')$. The map $T(S', p) \rightarrow T(S')$ forgets the point p . Elements of $T(S')$ represented by separating twists and bounding pair maps clearly lift to such elements of $T(S', p)$. We claim the kernel K_1 is generated by bounding pair maps. As in the proof of Lemma 8, the kernel of $M(S', p) \rightarrow M(S')$ consists of elements obtained by dragging p around a loop in S' . This is the image of the homomorphism $d : \pi_1(S', p) \rightarrow \pi_0 \text{Diff}^+(S', p)$ defined there. Generators for this kernel are obtained by dragging p around generators for $\pi_1(S', p)$, and we can choose embedded nonseparating curves for these generators, from the standard presentation of $\pi_1(S', p)$. Such drag maps are bounding pair maps, so the kernel of $M(S', p) \rightarrow M(S')$ is in fact contained in $T(S', p)$ and so coincides with K_1 . Thus K_1 is generated by bounding pair maps, as claimed. It follows that $T(S', p)$ is generated by separating twists and bounding pair maps if this is true for $T(S')$.

Now we proceed to the second short exact sequence. Separating twists and bounding pair maps in $T(S', p)$ lift to such maps in $T(S', P)$ by choosing the point q to be sufficiently close to p . Next we show that the kernel K_2 is generated by products of separating twists in $T(S', P)$. The kernel of $M(S', P) \rightarrow M(S, p)$ is formed by maps $d_\gamma = d(\gamma)$ resulting from dragging q around loops γ in $S' - p$, where now d is the boundary map $\pi_1(S' - p, q) \rightarrow \pi_0 \text{Diff}^+(S', P)$. To see how d_γ acts on $H_1(S', P)$, consider the short exact sequence

$$0 \rightarrow H_1(S') \rightarrow H_1(S', P) \rightarrow \tilde{H}_0(P) \rightarrow 0$$

which is invariant under $M(S', P)$. Clearly d_γ acts trivially on the image of $H_1(S')$ in $H_1(S', P)$ since d_γ is trivial in $M(S')$. Also d_γ acts trivially on $\tilde{H}_0(P)$ since it fixes P . The action of d_γ on $H_1(S', P)$ will then be trivial if and only if it acts trivially on an element of $H_1(S', P)$ that maps to a generator of $\tilde{H}_0(P)$. We can represent such an element by an arc ϵ joining p to q . It follows from the fact that

the diffeomorphisms d_γ act trivially on $H_1(S')$ and $\tilde{H}_0(P)$ that there is a homomorphism $\varphi : \pi_1(S' - p, q) \rightarrow H_1(S') \subset H_1(S', P)$ such that $d_\gamma[\epsilon] = [\epsilon] + \varphi(\gamma)$ for all γ . We claim that $\varphi(\gamma) = [\gamma]$, so that φ is the abelianization map $\pi_1(S' - p, q) \rightarrow H_1(S' - p) = H_1(S')$. Since φ is a homomorphism it suffices to check that $d_\gamma[\epsilon] = [\epsilon] + [\gamma]$ for γ ranging over a set of generators for $\pi_1(S' - p, q)$. As generators we can choose embedded loops γ disjoint from ϵ , and for such loops γ the formula $d_\gamma[\epsilon] = [\epsilon] + [\gamma]$ obviously holds.

From the formula $d_\gamma[\epsilon] = [\epsilon] + [\gamma]$ we see that d_γ acts trivially on $H_1(S', P)$ if and only if γ lies in the commutator subgroup of $\pi_1(S' - p)$. It is a general fact that the commutator subgroup of a group is normally generated by commutators of generators of the group. For $\pi_1(S' - p)$ we choose generators coming from representing $S' - p$ as a punctured $4(g - 1)$ -gon with opposite edges identified (this is not the standard identification!). These generators are nonseparating curves, any two of which intersect transversely in one point, so their commutator is represented by a curve γ bounding a genus 1 subsurface of $S' - p$. The map d_γ is then the composition of a twist along a parallel copy of γ and an inverse twist along another parallel copy of γ . This shows that K_2 is generated by products of separating twists. All such separating twists lie in $T(S', P)$ since for a twist along a separating curve γ that does not separate p and q a basis for $H_1(S, P)$ can be chosen disjoint from γ , while if γ does separate p and q then composing the twist along γ with d_γ or d_γ^{-1} converts the twist along γ to a twist along a γ that does not separate p and q , and $d_\gamma \in T(S', P)$.


In summary, we have shown the inductive step:

Proposition 10. *If $T(S')$ is generated by separating twists and bounding pair maps, then so are $T(S', p)$, $T(S', P)$, and $T(S)$ in turn, in the last case assuming that the genus of S is at least 3.*

The next two propositions, due to Johnson [9], justify two supplementary statements made earlier.

Proposition 11. *Every Dehn twist about a separating curve in a surface of genus at least 3 is isotopic to a product of bounding pair maps.*

Proof. First consider the case that S has genus 3. A nontrivial separating curve c in S then splits S into a punctured torus and a punctured genus 2 surface, and the latter surface can be further decomposed as the union of a 4-punctured sphere and a pair of pants, as in the figure below.



The lantern relation [3, p. 187] [9, §IV] [5, §5.1] gives

$$T_x T_y T_z = T_a T_b T_c T_d$$


Since each of T_a , T_b , T_c , and T_d commutes with all seven Dehn twists in the relation, we can rewrite this relation as follows:

$$(T_x T_a^{-1})(T_y T_b^{-1})(T_z T_d^{-1}) = T_c.$$

In other words the Dehn twist about the separating curve c is the product of three bounding pair maps. This takes care of the genus 3 case. The general case is obtained from this by attaching the appropriate number of handles to the punctured torus and the pair of pants. \square

Proposition 12. *Every bounding pair map of a surface S is isotopic to a product of bounding pair maps associated to pairs of curves bounding genus 1 subsurfaces of S .*


Proof. A bounding pair map is obtained by twisting a subsurface bounded by two curves through 360 degrees. If this subsurface has genus n , it can be decomposed into n subsurfaces of genus 1, each bounded by two curves, as indicated in the figure below, and the twist of the genus n subsurface is isotopic to the composition of n twists of the genus 1 subsurfaces. \square



5. STARTING THE INDUCTION: GENUS 2

Now we restrict to the case that S has genus 2, where we want to show that $T(S)$ is generated by separating twists. The complex $\mathcal{C}_x(S)$ is 0-dimensional in this case, so is of little help. Instead we use $\mathcal{B}_x(S)$, which is 1-dimensional and contractible, hence a tree. From the elementary theory of groups acting on trees, it will suffice to show that the quotient $\mathcal{B}_x(S)/T(S)$ is also a tree, and that the stabilizers of vertices in $\mathcal{B}_x(S)$ are generated by separating twists.


It will be helpful to know exactly what an edge of $\mathcal{B}_x(S)$ looks like. Such an edge corresponds to a multicurve separating S into two components, so these components must be pairs of pants, with the multicurve consisting of three nonseparating curves a, b, c as in the figure below.



Since the multicurve is reduced, two of the three curves, say a and b , will be oriented consistently with an orientation of either of the pairs of pants, and the third curve c will be oppositely oriented. As we move across the edge we transfer weights from a and b to c , or vice versa. Transferring weights from a and b to c decreases the value of W . At the end of the edge with larger W -value the weighted multicurve is $pa + qb$, with $p \geq q$ say, and then at the other end the weighted multicurve is $(p - q)a + qc$, so we subtract q from the weights on a and b and add q to the weight on c . The value of W decreases from $p + q$ to p . (Thus for a sequence of edges along which W decreases, the pairs of weights are changing according to the Euclidean algorithm of repeatedly subtracting the smaller of two numbers from the larger.)

We claim that from a given vertex $pa + qb$ with $W > 1$ (so both p and q are greater than 0) all the edges of $\mathcal{B}_x(S)$ leading to vertices with smaller W -value are equivalent under the action of the stabilizer of the vertex in $T(S)$. This implies that each component of $\mathcal{B}_x(S)/T(S)$ is a tree since there is a well-defined flow on it decreasing the values of W monotonically until they reach the value 1 at a vertex represented by a single curve. Since $\mathcal{B}_x(S)$ is connected, so is its quotient $\mathcal{B}_x(S)/T(S)$, so the quotient must then be a tree, with a single vertex where $W = 1$. In particular, this shows that $T(S)$ acts transitively on oriented curves in a given homology class, just as in higher genus as we noted in the remarks following Lemma 9.

To verify the claim, consider two edges of $\mathcal{B}_x(S)$ leading downward from this vertex. These two edges correspond to two different choices for the c curve, in the notation above. Both choices lie in the complement of the a and b curves, a 4-punctured sphere. Isotopy classes of nontrivial curves in a 4-punctured sphere are classified by their slope, an element of $\mathbb{Q} \cup \{\infty\}$. Let us choose coordinates on the 4-punctured sphere so that one choice of c has slope $0/1$ and a separating curve d as in the preceding figure has slope $1/0$. The mapping class group of a 4-punctured sphere fixing each of the punctures can be identified with the subgroup G of $PSL(2, \mathbb{Z})$ represented by matrices congruent to the identity mod 2. The action of G on slopes has three orbits, the slopes whose numerators and denominators are congruent to those of $0/1$, $1/0$, or $1/1$ mod 2. Topologically, these three classes are distinguished by how the corresponding curves separate the 4 punctures into pairs. In particular, the slopes congruent to $1/0$ correspond to separating curves on S , such as the curve d . Whether the value of W decreases or increases for a particular choice for c depends only on how c separates the 4 punctures. For c of slope $0/1$ as in the figure the value of W decreases, so the slopes congruent to $0/1$ correspond to the curves c for which W decreases. These slopes form the vertices of a tree T that can be visualized by superimposing it on the Farey diagram, as in the figure below.



Two vertices of T are joined by a sequence of edges of T , and the curves corresponding to vertices along this path are related each to the next by Dehn twists along curves of slope congruent to $1/0$ mod 2. For example, a twist along a curve of slope $1/0$ takes slope $0/1$ to slope $2/1$. Slopes congruent to $1/0$ mod 2 give separating curves in S , so

the path in T gives a product of separating twists on S taking the first choice of c to the second choice, verifying the claim.

Next we check that the stabilizer of a vertex of $\mathcal{B}_x(S)$ corresponding to a vertex $pa + qb$ with $p, q > 0$ is generated by separating twists. We have just seen that products of separating twists in the stabilizer act transitively on vertices of T , so, modulo such products, we can assume the element of the stabilizer fixes the slope 0/1 curve c . But the stabilizer of the multicurve $a \cup b \cup c$ in $T(S)$ is trivial, since an element in this stabilizer would have to be a product of twists along a , b , and c , and it is easy to see that such a product acts nontrivially on homology unless it is the trivial product. (Consider the action on curves that intersect two of a, b, c transversely in one point and are disjoint from the third.) Thus the stabilizer of $pa + qb$ in $T(S)$ is generated by separating twists when both p and q are positive.

There remains the stabilizer of a vertex corresponding to a single curve a . This situation was analyzed in the previous section for arbitrary genus, where we showed that the stabilizer of a curve in genus g is generated by twists and bounding pair maps if this is true for the full Torelli group in one lower genus. In the present situation the Torelli group for genus 1 is trivial, so the stabilizer of a curve in genus 2 is generated by separating twists since there are no bounding pair maps until genus 3.

This finishes the proof that $T(S)$ is generated by separating twists in genus 2, and hence also the proof of the Birman–Powell theorem.

The structure of the Torelli group in genus 2. It is not hard to extend the preceding arguments to see that $T(S)$ is a nonfinitely generated free group in genus 2. For the action of $T(S)$ on the tree $\mathcal{B}_x(S)$ the edge stabilizers are trivial as we observed above, and the quotient $\mathcal{B}_x(S)/T(S)$ is a tree, so $T(S)$ is a free product of vertex stabilizers, with one factor for each vertex of $\mathcal{B}_x(S)/T(S)$. For a vertex $pa + qb$ of $\mathcal{B}_x(S)$ with $p, q > 0$ the stabilizer acts freely on the tree T since we saw that no vertices can be fixed points, and no edge can be inverted since elements of the group G cannot interchange slopes congruent to 1/0 and 1/1 mod 2. The stabilizer group acts transitively on vertices of T , so the stabilizer is the fundamental group of the orbit space of the action, an infinite wedge of circles since vertices of T have infinite valence. Thus the stabilizer of $pa + qb$ is a free group on an infinite number of generators.

The other case is the stabilizer of a vertex that is a single curve. From the discussion in the preceding section this is isomorphic to the kernel K_2 of the map $T(S, P) \rightarrow T(S, p)$. This is a subgroup of the kernel of

the map $M(S, P) \rightarrow M(S, p)$ which is $\pi_1(S - p)$ from the long exact sequence of homotopy groups for the fibration $\text{Diff}^+(M, p) \rightarrow M - p$ obtained by evaluating diffeomorphisms at q , with fiber $\text{Diff}^+(M, P)$. The group $\pi_1(M - p)$ is free so the subgroup K_2 is free as well. It is in fact the commutator subgroup, as our analysis showed, so it is nonfinitely generated.

Since $\mathcal{B}_x(S)/T(S)$ is an infinite tree, we see that $T(S)$ is a free product of infinitely many stabilizers, each of which is a nonfinitely generated free group, so $T(S)$ is a nonfinitely generated free group itself. Mess [13] gives the following more precise description of the infinite generating set. Each nontrivial separating curve in S induces a splitting of $H_1(S)$ into two symplectic subspaces, namely, the subspaces consisting of the elements represented by 1-cycles supported entirely on one side of the separating curve or the other. Mess proved that $T(S)$ has a free generating set where there is one generator for each such symplectic splitting of $H_1(S)$ (this description can in fact be deduced by sharpening the argument given above, as in [1, Section 7]). It is not true, however, that if we make an arbitrary choice of Dehn twist for each symplectic splitting, then we obtain a generating set. As such, it is still an open problem to turn Mess's description of the generating set into an explicit generating set.

Final remark. It is tempting to try to prove the Birman–Powell theorem in genus 2 using the same inductive step we used in higher genus. Indeed, the genus one Torelli group is trivial, so by the Birman exact sequence it would suffice to show that $\mathcal{C}_x(S)$, or some variant, is connected in genus 2.

We have already mentioned that $\mathcal{C}_x(S)$ has no edges in genus 2. In particular, the complex is not connected. One might try to repair this by enlarging $\mathcal{C}_x(S)$ to a complex with edges joining pairs of vertices corresponding to curves that are not disjoint but intersect in the minimum number of points, namely 4, such as the curves c and $T_d(c)$ where c and d are the curves shown in the figure at the beginning of this section. This does not work, however. The curves c and $T_d(c)$ are joined by an edge path in $\mathcal{B}_x(S)$ of length 2 with $a + b$ as the intermediate vertex, so the values of W along this edge path lie between 1 and 2. But there are pairs of vertices of $\mathcal{C}_x(S)$ for which the values of W along the path in $\mathcal{B}_x(S)$ joining the two vertices must exceed any preassigned number n , since one can start with a vertex of $\mathcal{B}_x(S)$ where W has a value larger than n and then follow two different paths from this vertex along which W decreases monotonically until one reaches a pair of vertices in $\mathcal{C}_x(S)$ with $W = 1$. Since $\mathcal{B}_x(S)$ is a tree, these two vertices

cannot be joined by any other path along which W has the maximum value 2, so these two vertices cannot be in the same path component of the proposed enlargement of $\mathcal{C}_x(S)$. This argument shows moreover that $\mathcal{C}_x(S)$ cannot be made connected by adding only a finite number of types of edges.

REFERENCES

- [1] Mladen Bestvina, Kai-Uwe Bux, and Dan Margalit. The dimension of the Torelli group. *J. Amer. Math. Soc.*, 23(1):61–105, 2010.
- [2] Joan S. Birman. On Siegel’s modular group. *Math. Ann.*, 191:59–68, 1971.
- [3] Max Dehn. Die Gruppe der Abbildungsklassen. *Acta Math.*, 69(1):135–206, 1938. Das arithmetische Feld auf Flächen.
- [4] 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.
- [5] Benson Farb and Dan Margalit. *A primer on mapping class groups*. Princeton University Press, 2011.
- [6] Dennis Johnson. An abelian quotient of the mapping class group \mathcal{I}_g . *Math. Ann.*, 249(3):225–242, 1980.
- [7] 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.
- [8] 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.
- [9] Dennis L. Johnson. Homeomorphisms of a surface which act trivially on homology. *Proc. Amer. Math. Soc.*, 75(1):119–125, 1979.
- [10] Helmut Klingens. Charakterisierung der Siegelschen Modulgruppe durch ein endliches System definierender Relationen. *Math. Ann.*, 144:64–82, 1961.
- [11] 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.
- [12] Darryl McCullough and Andy Miller. The genus 2 Torelli group is not finitely generated. *Topology Appl.*, 22(1):43–49, 1986.
- [13] Geoffrey Mess. The Torelli groups for genus 2 and 3 surfaces. *Topology*, 31(4):775–790, 1992.
- [14] Jerome Powell. Two theorems on the mapping class group of a surface. *Proc. Amer. Math. Soc.*, 68(3):347–350, 1978.
- [15] Andrew Putman. Cutting and pasting in the Torelli groups. *Geom. Topol.*, 11:829–865 (electronic), 2007.
- [16] Andrew Putman. A note on the connectivity of certain complexes associated to surfaces. *Enseign. Math. (2)*, 54(3-4):287–301, 2008.
- [17] Andrew Putman. Small generating sets for the Torelli group. 2011. arXiv:1106.3294.
- [18] Barbara van den Berg. On the abelianization of the Torelli group. *Ph.D. thesis, University of Utrecht*, 2003.