Lectures on the Torelli group

Andrew Putman

Department of Mathematics, Rice University MS-136, 6100 Main St., Houston, TX 77005

E-mail address: andyp@rice.edu

2010 *Mathematics Subject Classification.* Primary 20F05; Secondary: 20F38, 57M07, 57N05

Key words and phrases. Torelli group, mapping class group, Johnson homomorphism, Johnson filtration

The author was supported in part by NSF grant DMS-1255350 and the Alfred P. Sloan Foundation.

Contents

Part 1

Foundational Material

Introduction to Part 1

Part 1 is devoted to fundamental topics that are used throughout the remainder of the book. Chapter 1 is an introduction to the mapping class group, Chapter 2 discusses the symplectic representation of the mapping class group obtained from its action on the first homology group of the surface, and Chapter 3 defines the Torelli group and discusses its basic properties.

CHAPTER 1

The mapping class group

This chapter introduces the mapping class group of a surface. Though we will discuss all the results needed elsewhere in this book, we will omit some of the lengthier proofs. There are many fine sources that discuss this material. We learned much of it from Ivanov's survey [**Iva02**] and Farb–Margalit's book [**FM12**], and we refer the interested reader to these sources for the missing proofs (and much more). Many readers will probably skip this chapter, though we recommend skimming it to learn our notational conventions.

1.1. Definitions and basic examples

We start by defining the mapping class group.

Surfaces and the mapping class group. In this book, a *surface* will always mean a compact oriented surface with boundary. If Σ is a surface, then the *mapping class group* of Σ, denoted $\text{Mod}(\Sigma)$, is the group of isotopy classes of orientation-preserving diffeomorphisms of Σ that restrict to the identity on the boundary. We emphasize that the isotopies must themselves fix the boundary pointwise. We will often denote a compact oriented genus g surface with b boundary components by Σ_g^b and its mapping class group by Mod_g^b; the b will sometimes be omitted when it vanishes.

Homotopies vs isotopies. A fundamental result of Baer [**Bae27, Bae28**] says that two diffeomorphisms of a closed orientable surface which are homotopic are also isotopic. This was later extended by Epstein [**Eps66**] to diffeomorphisms of orientable surfaces with boundary that fix the boundary pointwise (Epstein worked in the PL category, but it is easy to modify his proof to work in the smooth category). We therefore do not need to worry about the distinction between homotopies and isotopies. In fact, for much of this book it would be reasonable to simply define the mapping class group as the group of homotopy classes of orientation-preserving diffeomorphisms. However, this would cause some small technical difficulties later when we use mapping classes to glue 3-manifolds together along their boundaries; the result would not be obviously well-defined.

Discs and spheres. We now turn to the mapping class groups of low-complexity surfaces. We start with the disc $\mathbb{D}^2 = \Sigma_0^1$.

PROPOSITION 1.1. $\text{Mod}(\mathbb{D}^2) = 1$.

PROOF. Consider an orientation-preserving diffeomorphism $F : \mathbb{D}^2 \to \mathbb{D}^2$ such that $F|_{\partial \mathbb{D}^2} = \text{id}$. Regarding $\mathbb{D}^2$ as a subset of $\mathbb{C}$, we can homotope F to the identity via the straight-line homotopy

$$F_t : \mathbb{D}^2 \to \mathbb{D}^2$$

$$F_t(x) = (1-t)F(x) + tx \qquad \square$$

REMARK 1.2. Observe that in the proof of Proposition 1.1, the homotopy we wrote down is not necessarily an isotopy, so we are silently appealing to the aforementioned theorem of Epstein. With a bit more care one can directly produce an isotopy; see the discussion in [**FM12**]. We will ignore the distinction between homotopies and isotopies in many proofs in this section.

A similar result holds for the sphere $S^2 = \Sigma_0$.

PROPOSITION 1.3. *$\mathrm{Mod}(S^2) = 1$.*

PROOF. Consider an orientation-preserving diffeomorphism $F : S^2 \to S^2$. Fixing a basepoint $p_0 \in S^2$, we homotope F such that $F(p_0) = p_0$ by postcomposing F with rotations of S^2. Let $U \subset S^2$ be the open hemisphere centered at p_0. Both U and $F(U)$ are tubular neighborhoods of p_0, so by the usual uniqueness up to isotopy of tubular neighborhoods (see, e.g., [**Hir94**, Theorem 4.5.3]; this is where we use the fact that F is orientation-preserving) we can isotope F such that $F|_U = \mathrm{id}$. Set $D = S^2 \backslash U$, so $D \cong \mathbb{D}^2$. The map F restricts to a diffeomorphism $F|_D : D \to D$ that fixes ∂D pointwise. Using Proposition 1.1, we can therefore homotope F to the identity. $\square$

The annulus. We now turn to the annulus $\mathbb{A} = \Sigma_0^2$, which provides us with our first example of a nontrivial mapping class.

PROPOSITION 1.4. *$\mathrm{Mod}(\mathbb{A}) = \mathbb{Z}$.*

PROOF. We first define a homomorphism $\psi : \mathrm{Mod}(\mathbb{A}) \to \mathbb{Z}$. Let $\pi : \widetilde{\mathbb{A}} \to \mathbb{A}$ be the universal cover. We will identify $\widetilde{\mathbb{A}}$ with $\{z \in \mathbb{C} \mid 0 \leqslant \mathrm{Im}(z) \leqslant 1\}$; the deck group $\mathbb{Z}$ acts by horizontal translations $z \mapsto z + n$. Define

$$B = \{x \mid x \in \mathbb{R}\} \subset \widetilde{\mathbb{A}} \quad \text{and} \quad T = \{x + i \mid x \in \mathbb{R}\} \subset \widetilde{\mathbb{A}}.$$

Consider $f \in \mathrm{Mod}(\mathbb{A})$ which is represented by an orientation-preserving diffeomorphism $F : \mathbb{A} \to \mathbb{A}$ with $F|_{\partial\mathbb{A}} = \mathrm{id}$. We can uniquely lift F to a diffeomorphism $\widetilde{F} : \widetilde{\mathbb{A}} \to \widetilde{\mathbb{A}}$ satisfying $F(0) = 0$. This latter condition implies that $\widetilde{F}|_B = \mathrm{id}$. However, we do not necessarily have $\widetilde{F}|_T = \mathrm{id}$; instead, $\widetilde{F}(x + i) = (x + n_F) + i$ for some $n_F \in \mathbb{Z}$. The homotopy lifting property implies that n_F is unchanged by homotopies of F that fix $\partial\mathbb{A}$ pointwise. We can therefore define $\psi(f) = n_F$. It is clear that ψ is a homomorphism.

To see that ψ is injective, consider $h \in \mathrm{Mod}(\mathbb{A})$ such that $\psi(h) = 0$. Letting $H : \mathbb{A} \to \mathbb{A}$ be a representative diffeomorphism, we can lift H to a diffeomorphism $\widetilde{H} : \widetilde{\mathbb{A}} \to \widetilde{\mathbb{A}}$ satisfying $\widetilde{H}|_{\partial\widetilde{\mathbb{A}}} = \mathrm{id}$. The straight-line homotopy

$$\widetilde{H}_t : \widetilde{\mathbb{A}} \to \widetilde{\mathbb{A}}$$
$$\widetilde{H}_t(z) = (1-t)\widetilde{H}(z) + tz$$

from $\widetilde{H}$ to id commutes with the deck group, and therefore projects to a homotopy from H to id. This projected homotopy fixes $\partial\mathbb{A}$ pointwise (this is where we use the fact that $\psi(h) = 0$), so we conclude that $h = 1$.

To see that ψ is surjective, consider some $n \in \mathbb{Z}$. Defining

$$\widetilde{F}_n : \widetilde{\mathbb{A}} \to \widetilde{\mathbb{A}}$$

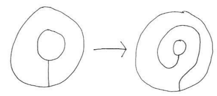

FIGURE 1.1. $\mathrm{Mod}(\mathbb{A})$ is generated by the mapping class that holds the outer boundary of $\mathbb{A}$ fixed while rotating the inner boundary by 2π; as shown here, this causes an arc connecting the two boundary components to acquire a segment going around the annulus.

$$\widetilde{F}_n(z) = z + n\,\mathrm{Im}(z),$$

the diffeomorphism $\widetilde{F}_n$ commutes with the deck group and therefore projects to a diffeomorphism $F_n : \mathbb{A} \to \mathbb{A}$. By construction we have $F_n|_{\partial\mathbb{A}} = \mathrm{id}$, so F_n defines a mapping class $f_n \in \mathrm{Mod}(\mathbb{A})$ which satisfies $\psi(f_n) = n$. □

REMARK 1.5. The generator for $\mathrm{Mod}(\mathbb{A}) \cong \mathbb{Z}$ constructed in the above proof is illustrated in Figure 1.1.

The algebraic intersection pairing. Below we will study the mapping class group of the 2-torus $\mathbb{T} = \Sigma_1$. To do this, we will need the *algebraic intersection* pairing. For a surface Σ, this is a $\mathbb{Z}$-valued bilinear form $\hat{i}(\cdot,\cdot)$ on $\mathrm{H}_1(\Sigma;\mathbb{Z})$ which satisfies the following two properties (see [**Bre97**, §VI.11] for more details).

- It is *alternating* in the sense that $\hat{i}(h_1,h_2) = -\hat{i}(h_2,h_1)$ for all $h_1,h_2 \in \mathrm{H}_1(\Sigma;\mathbb{Z})$. This implies in particular that $\hat{i}(h,h) = 0$ for all $h \in \mathrm{H}_1(\Sigma;\mathbb{Z})$.
- For $h_1,h_2 \in \mathrm{H}_1(\Sigma;\mathbb{Z})$, the number $\hat{i}(h_1,h_2) \in \mathbb{Z}$ can be calculated as follows. Choose cycles c_1 and c_2 representing h_1 and h_2, respectively, such that the c_i intersect transversely. We then have
$$\hat{i}(h_1,h_2) = \sum_{p \in c_1 \cap c_2} \pm 1,$$
where the sign ±1 is the sign of the intersection at p.

We will have much more to say about $\hat{i}(\cdot,\cdot)$ in Chapter 2.

The torus. The group $\mathrm{Mod}(\mathbb{T})$ is richer than any of the other mapping class groups considered so far, and its study will lead us to the main topic of this book. Observe that the group $\mathrm{Mod}(\mathbb{T})$ acts on $\mathrm{H}_1(\mathbb{T};\mathbb{Z}) \cong \mathbb{Z}^2$. This action induces a homomorphism $\mathrm{Mod}(\mathbb{T}) \to \mathrm{Aut}(\mathbb{Z}^2) \cong \mathrm{GL}_2(\mathbb{Z})$. However, this map is *not* surjective. The issue is that the image of this map preserves the algebraic intersection pairing $\hat{i}(\cdot,\cdot)$ on $\mathrm{H}_1(\mathbb{T};\mathbb{Z})$. Let α and β be the two oriented curves on $\mathbb{T}$ depicted in Figure 1.2 and let a and b be their homology classes, so $\hat{i}(a,b) = 1$. The homology classes a and b form a basis for $\mathrm{H}_1(\mathbb{T};\mathbb{Z})$. For $f \in \mathrm{Mod}(\mathbb{T})$, write $f(a) = c_1 a + c_2 b$ and $f(b) = d_1 a + d_2 b$ with $c_1, c_2, d_1, d_2 \in \mathbb{Z}$. We then have

$$1 = \hat{i}(a,b) = \hat{i}(f(a),f(b)) = \hat{i}(c_1 a + c_2 b, d_1 a + d_2 b) = c_1 d_2 - c_2 d_1;$$

the minus sign appears because the algebraic intersection pairing is alternating. The expression $c_1 d_2 - c_2 d_1$ is the determinant of the action of ψ on $\mathrm{H}_1(\mathbb{T};\mathbb{Z})$. The upshot is that the action of $\mathrm{Mod}(\mathbb{T}^2)$ on $\mathrm{H}_1(\mathbb{T}^2;\mathbb{Z})$ yields a homomorphism $\mathrm{Mod}(\mathbb{T}^2) \to \mathrm{SL}_2(\mathbb{Z})$. The following proposition says that this is an isomorphism.

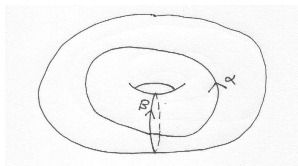

FIGURE 1.2. The curves α and β on the torus $\mathbb{T}^2$ whose homology classes a and b generate $\mathrm{H}_1(\mathbb{T}^2;\mathbb{Z})$ and satisfy $\hat{i}(a,b) = 1$.

PROPOSITION 1.6. *The map* $\mathrm{Mod}(\mathbb{T}^2) \to \mathrm{SL}_2(\mathbb{Z})$ *obtained from the action of* $\mathrm{Mod}(\mathbb{T}^2)$ *on* $\mathrm{H}_2(\mathbb{T}^2;\mathbb{Z})$ *is an isomorphism.*

PROOF. Let $\rho : \mathrm{Mod}(\mathbb{T}^2) \to \mathrm{SL}_2(\mathbb{Z})$ be the map in question. Regard $\mathbb{T}^2$ as the quotient of $\mathbb{R}^2$ by $\mathbb{Z}^2$, and let $\pi : \mathbb{R}^2 \to \mathbb{T}^2$ be the projection.

To see that ρ is surjective, consider $M \in \mathrm{SL}_2(\mathbb{Z})$. The action of M on $\mathbb{R}^2$ gives a diffeomorphism $\widetilde{F}_M : \mathbb{R}^2 \to \mathbb{R}^2$ that projects to a diffeomorphism $F_M : \mathbb{T}^2 \to \mathbb{T}^2$. The resulting mapping class $f_M \in \mathrm{Mod}(\mathbb{T}^2)$ clearly satisfies $\rho(f_M) = M$.

To see that ρ is injective, consider $f \in \ker(\rho)$. Choose a diffeomorphism $F : \mathbb{T}^2 \to \mathbb{T}^2$ representing f. Letting $p_0 = \pi(0)$, we can homotope F such that $F(p_0) = p_0$. Lift F to a diffeomorphism $\widetilde{F} : \mathbb{R}^2 \to \mathbb{R}^2$ satisfying $\widetilde{F}(0) = 0$. Since $f \in \ker(\rho)$, it follows that $\widetilde{F}(x+n, y+m) = \widetilde{F}(x,y) + (n,m)$ for all $(x,y) \in \mathbb{R}^2$ and $(n,m) \in \mathbb{Z}^2$. This implies that the straight-line homotopy

$$\widetilde{F}_t : \mathbb{R}^2 \to \mathbb{R}^2$$

$$\widetilde{F}_t(x,y) = (1-t)\widetilde{F}(x,y) + t(x,y)$$

projects to a homotopy from F to id, so $f = 1$. □

Higher genus. For genus at least 2, there is no simple description of the mapping class group analogous to Proposition 1.6. Indeed, while there is still a representation $\mathrm{Mod}(\Sigma) \to \mathrm{Aut}(\mathrm{H}_1(\Sigma;\mathbb{Z}))$, this representation is far from injective. Its kernel is known as the *Torelli group* and is the main subject of this book.

An abuse of notation. In the above proofs, we maintained the distinction between an element of the mapping class group and a diffeomorphism representing it. Continuing to do this would seriously complicate our notation, so as is traditional in the mapping class group literature we will cease to make this distinction (except in a few cases where this might lead to confusion).

1.2. Dehn twists

We can parlay the fact that $\mathrm{Mod}(\mathbb{A}) \cong \mathbb{Z}$ from Proposition 1.4 above into a construction of an important class of elements of $\mathrm{Mod}(\Sigma)$ for an arbitrary surface Σ.

Dehn twists. Consider a simple closed curve γ on Σ. Let A_γ be a closed tubular neighborhood of γ, so $A_\gamma \cong \mathbb{A}$. There is a natural map $\mathrm{Mod}(A_\gamma) \to \mathrm{Mod}(\Sigma)$ that extends a mapping class on A_γ to Σ by the identity (this works because we required that mapping classes act as the identity on the boundary). The image of a generator of $\mathrm{Mod}(A_\gamma) \cong \mathbb{Z}$ is a *Dehn twist*. Of course, there are two generators

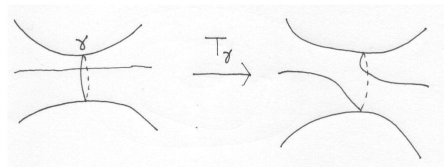

FIGURE 1.3. The effect of a Dehn twist T_γ.

of $\mathbb{Z}$; we will denote by $T_\gamma \in \mathrm{Mod}(\Sigma)$ the image of the generator of $\mathrm{Mod}(A_\gamma)$ that has the effect depicted in Figure 1.3 (the other generator goes to T_γ^{-1}; the mapping class T_γ is sometimes called a "right-handed Dehn twist" and the mapping class T_γ^{-1} is sometimes called a "left-handed Dehn twist").

Properties of Dehn twists. Here are several important properties of T_γ.

(1) The mapping class T_γ does not depend on the choice of N. This follows from the uniqueness up to isotopy of tubular neighborhoods; see, e.g., [**Hir94**, Theorem 4.5.3].
(2) If γ and γ' are isotopic simple closed curves, then $T_\gamma = T_{\gamma'}$. This is immediate once the previous property is established.
(3) If γ is not nullhomotopic, then T_γ is an infinite-order element of $\mathrm{Mod}(\Sigma)$; see [**FM12**, Chapter 3]. If γ is instead nullhomotopic, then $T_\gamma = 1$.

In light of the second property above, it makes sense to talk about the Dehn twist about an isotopy class of simple closed curves. In fact, for the most part in this book we will not distinguish between a simple closed curve and its isotopy class (this is similar to the fact that we will usually not distinguish between a mapping class and a diffeomorphism representing that mapping class).

The torus. Let α and β be the simple closed curves on the torus $\mathbb{T}^2$ depicted in Figure 1.2 and let a and b be their homology classes. Recalling from Proposition 1.6 that $\mathrm{Mod}(\mathbb{T}^2) \cong \mathrm{SL}_2(\mathbb{Z})$, it is easy to see that with respect to the basis $\{a, b\}$ for $\mathrm{H}_1(\mathbb{T}^2; \mathbb{Z})$, the Dehn twists T_α and T_β correspond to the matrices

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 1 & 0 \\ -1 & 1 \end{pmatrix},$$

respectively.

Generating the mapping class group. Dehn twists were discovered by Dehn [**Deh38**] in 1938, but were forgotten until they were rediscovered by Lickorish [**Lic64**] in 1964 (for a while, they were known as "Lickorish twists"). Their importance is underlined by the following theorem, which was also proven by both Dehn and Lickorish (at least for closed surfaces).

THEOREM 1.7. *If* Σ *is a surface, then* $\mathrm{Mod}(\Sigma)$ *is generated by the set of all Dehn twists.*

The proof of Theorem 1.7 is lengthy, so we will omit it (see [**FM12**] or [**Iva02**] for the details).

Finite generation. In fact, even more is true : the mapping class group is generated by *finitely many* Dehn twists. For closed surfaces Σ_g, this was proved

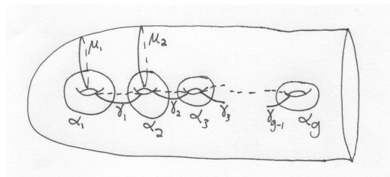

FIGURE 1.4. The Dehn twists about the curves $\alpha_1, \dots, \alpha_g$ and $\gamma_1, \dots, \gamma_{g-1}$ and μ_1, μ_2 generate Mod_g^b. The figure depicts Σ_g^1; to get the curves on Σ_g, glue a disc to the boundary component.

by Dehn [**Deh38**], who found a set of $2g(g-1)$ Dehn twists that generated Mod_g. Lickorish [**Lic64**] later proved that $3g-1$ Dehn twists suffice to generate Mod_g. The definitive result in this direction is due to Humphries [**Hum79**], who found a set of $2g+1$ Dehn twists that generate Mod_g and proved that no smaller set sufficed. One can show that the same result also holds for surfaces with one boundary component; in fact, we have the following.

THEOREM 1.8. *For $g \geqslant 2$ and $0 \leqslant b \leqslant 1$, the group Mod_g^b is generated by the $2g+1$ Dehn twists $T_{\alpha_1}, \dots, T_{\alpha_g}, T_{\gamma_1}, \dots, T_{\gamma_{g-1}}, T_{\mu_1}, T_{\mu_2}$ depicted in Figure 1.4.*

Again, the proof is lengthy and thus omitted; see [**FM12**] or [**Iva02**] for the details. Also see [**FM12**, §4.4.4] for an explicit set of Dehn twists that generate Mod_g^b for $b \geqslant 2$.

1.3. The classification of surfaces trick

In this book, the word *curve* will always mean the homotopy class of a curve. One of the most important techniques for studying the mapping class group is to utilize its action on the set of curves on the surface. This section is devoted to an important trick that we will call the *classification of surfaces trick* which elucidates this action.

Single curve. The proof of the following lemma is an easy example of the classification of surfaces trick.

LEMMA 1.9. *Let Σ be a surface and let α and α' be oriented simple closed curves on Σ. Assume that neither α nor α' separate the surface. Then there exists some $f \in \mathrm{Mod}(\Sigma)$ such that $f(\alpha) = \alpha'$.*

PROOF. Let $b \geqslant 0$ be the number of boundary components of Σ. Let Σ_α and $\Sigma_{\alpha'}$ be the surfaces that result from cutting Σ along α and α', respectively. Since neither α nor α' separate Σ, both Σ_α and $\Sigma_{\alpha'}$ are connected surfaces. It is also clear that $\chi(\Sigma_\alpha) = \chi(\Sigma_{\alpha'})$ and that both Σ_α and $\Sigma_{\alpha'}$ have $b+2$ boundary components. The classification of surfaces therefore says that they are diffeomorphic. Let ∂_1 and ∂_2 (resp. ∂_1' and ∂_2') be the boundary components of Σ_α (resp. $\Sigma_{\alpha'}$) coming from α (resp. α'). The orientations of α and α' induce orientations on the ∂_i and the ∂_i'; order them so that Σ lies to the left of ∂_1 and ∂_1' and to the right of ∂_2 and ∂_2'. We can then choose an orientation-preserving diffeomorphism $\phi : \Sigma_\alpha \to \Sigma_{\alpha'}$ such that $\phi(\partial_i) = \partial_i'$ (as oriented curves) for $i = 1, 2$ and such that ϕ matches up the boundary components of Σ_α and $\Sigma_{\alpha'}$ that come from Σ. Gluing the ∂_i and ∂_i'

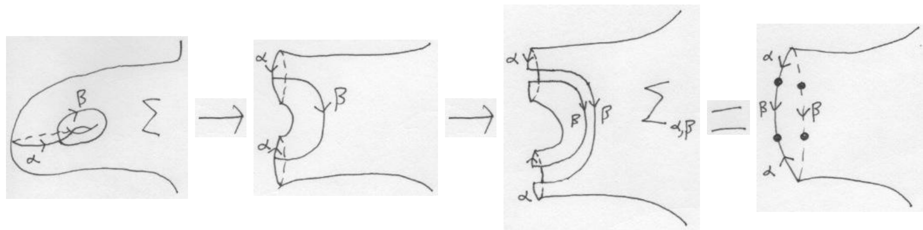

FIGURE 1.5. Cutting Σ first along α and then along β causes the genus to drop by 1 and the number of boundary components to increase by 1.

back together, we obtain a diffeomorphism $\psi : \Sigma \to \Sigma$ with $\psi|_{\partial\Sigma} = \text{id}$ satisfying $\psi(\alpha) = \alpha'$, as desired. □

This has the following corollary.

COROLLARY 1.10. *Let Σ be a surface and let α and α' be nonseparating simple closed curves on Σ. Then T_α and $T_{\alpha'}$ are conjugate in* $\text{Mod}(\Sigma)$.

PROOF. Orienting the α arbitrarily, Lemma 1.9 says that there exists some $f \in \text{Mod}(\Sigma)$ such that $f(\alpha) = \alpha'$. We then have the following calculation; the first equality is an easy exercise.

$$fT_\alpha f^{-1} = T_{f(\alpha)} = T_{\alpha'}$$ □

Handles. Here is another example of the classification of surfaces trick.

LEMMA 1.11. *Let Σ be a surface and let $\{\alpha, \beta\}$ and $\{\alpha', \beta'\}$ be collections of oriented simple closed curves on Σ. Assume that α and β intersect once with a positive sign. Similarly, assume that α' and β' intersect once with a positive sign. Then there exists some $f \in \text{Mod}(\Sigma)$ such that $f(\alpha) = \alpha'$ and $f(\beta) = \beta'$.*

PROOF. Assume that $\Sigma \cong \Sigma_g^b$. Define $\Sigma_{\alpha,\beta}$ and $\Sigma_{\alpha',\beta'}$ be the surfaces that result from cutting Σ along $\alpha \cup \beta$ and $\alpha' \cup \beta'$, respectively. As is shown in Figure 1.5, we have $\Sigma_{\alpha,\beta} \cong \Sigma_{g-1}^{b+1}$. This figure also shows that the boundary component ∂ of $\Sigma_{\alpha,\beta}$ coming from $\alpha \cup \beta$ can be divided into four oriented arcs, two of which glue up to form α and two of which glue up to form β. A similar thing is true for $\Sigma_{\alpha',\beta'}$ and its new boundary component ∂'. Using the classification of surfaces, we can find an orientation-preserving homeomorphism $\phi : \Sigma_{\alpha,\beta} \to \Sigma_{\alpha',\beta'}$ that matches up the boundary components coming from Σ and that takes ∂ to ∂'. Moreover, we can choose ϕ such that it respects the division of ∂ and ∂' into oriented arcs and takes the arcs corresponding to α and β to the arcs corresponding to α' and β', respectively (this is where we use the fact that the intersections have positive sign; otherwise, we might not be able to match up the orientations on these arcs). Gluing ∂ and ∂' back together, we obtain a diffeomorphism $\psi : \Sigma \to \Sigma$ with $\psi|_{\partial\Sigma} = \text{id}$ satisfying $\psi(\alpha) = \alpha'$ and $\psi(\beta) = \beta'$, as desired. □

Other examples of trick. The trick in the proofs of Lemmas 1.9–1.11 can be used in a wide variety of situations to show that the mapping class group acts transitively on collections of submanifolds of a surface that "cut the surface up in

FIGURE 1.6. Examples of configurations in Lemmas 1.12–1.14

the same way". Here are some other examples of it in action. We leave their proofs as exercises. Figure 1.6 gives examples of the configurations in these lemmas.

LEMMA 1.12. *Let Σ be a surface and let $\{\alpha_1, \ldots, \alpha_k\}$ and $\{\alpha'_1, \ldots, \alpha'_k\}$ be collections of oriented simple closed curves on Σ. Assume that the α_i are pairwise disjoint and that $\alpha_1 \cup \cdots \cup \alpha_k$ does not disconnect Σ. Similarly, assume that the α'_i are pairwise disjoint and that $\alpha'_1 \cup \cdots \cup \alpha'_k$ does not disconnect Σ. Then there exists some $f \in \text{Mod}(\Sigma)$ such that $f(\alpha_i) = \alpha'_i$ for $1 \leqslant i \leqslant k$.*

LEMMA 1.13. *Let Σ be a closed surface and let γ and γ' be simple closed curves on Σ. Assume that γ separates Σ into two subsurfaces S_1 and S_2 and that γ' separates Σ into two subsurfaces S'_1 and S'_2. Furthermore, assume that $S_i \cong S'_i$ for $i = 1, 2$. Then there exists some $f \in \text{Mod}(\Sigma)$ such that $f(\gamma) = \gamma'$.*

LEMMA 1.14. *Let Σ be a closed surface and let $\{\gamma_1, \gamma_2\}$ and $\{\gamma'_1, \gamma'_2\}$ be collections of simple closed curves on Σ. Assume that neither γ_1 nor γ_2 separate Σ but that $\gamma_1 \cup \gamma_2$ separates Σ into two subsurfaces S_1 and S_2. Similarly, assume that neither γ'_1 nor γ'_2 separate Σ but that $\gamma'_1 \cup \gamma'_2$ separates Σ into two subsurfaces S'_1 and S'_2. Furthermore, assume that $S_i \cong S'_i$ for $i = 1, 2$. Then there exists some $f \in \text{Mod}(\Sigma)$ such that $f(\gamma_i) = \gamma'_i$ for $i = 1, 2$.*

REMARK 1.15. In Lemma 1.14, the pairs $\{\gamma_1, \gamma_2\}$ and $\{\gamma'_1, \gamma'_2\}$ form what we will call *bounding pairs* in Chapter 3.

REMARK 1.16. There are also versions of Lemmas 1.13–1.14 for surfaces with boundary and for oriented curves, but they require small tweaks in their statements; we invite the reader to figure out the appropriate generalizations.

For the reader who has not seen this idea before, we recommend perusing [**FM12**, §1.1.3], which contains many examples of it (we remark that [**FM12**] calls it the "change of coordinates principle").

1.4. The Birman exact sequence and curve stabilizers

The Birman exact sequence is a basic tool that relates the mapping class groups of surfaces with differing numbers of boundary components. We will use it to understand the stabilizers in the mapping class group of nonseparating simple closed curves.

Statement. The form of the Birman exact sequence we will use was first proved by Johnson [**Joh83**]. It is a variant on a theorem of Birman [**Bir69**] which dealt with punctured surfaces instead of surfaces with boundary.

THEOREM 1.17. *Let Σ be a surface such that $\Sigma \ncong \Sigma_1^1$ and let β be a boundary component of Σ. Define $\widehat{\Sigma}$ to be the result of gluing a disc to β. Then there is a*

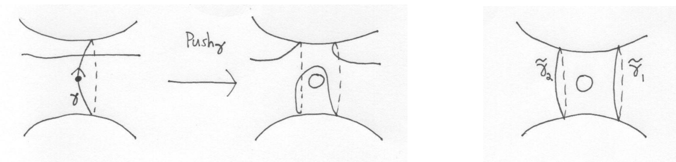

FIGURE 1.7. The effect of pushing a boundary component β of Σ about a simple closed curve $\gamma \in \pi_1(\widehat{\Sigma})$ is $T_{\tilde{\gamma}_1} T_{\tilde{\gamma}_2}^{-1}$.

short exact sequence

$$1 \longrightarrow \pi_1(U\widehat{\Sigma}) \longrightarrow \mathrm{Mod}(\Sigma) \longrightarrow \mathrm{Mod}(\widehat{\Sigma}) \longrightarrow 1,$$

where $U\widehat{\Sigma}$ is the unit tangent bundle of $\widehat{\Sigma}$.

The map $\mathrm{Mod}(\Sigma) \to \mathrm{Mod}(\widehat{\Sigma})$ in Theorem 1.17 is the map which extends mapping classes over the glued-in disc by the identity. The mapping classes in the kernel $\pi_1(U\widehat{\Sigma}) \subset \mathrm{Mod}(\Sigma)$ "drag" the boundary component β around the surface while allowing it to rotate; the loop around the fiber in $U\widehat{\Sigma}$ corresponds to T_β. See [**FM12**, §4.2.5] for a proof of Theorem 1.17, which we omit.

One-holed torus. The condition $\Sigma \not\cong \Sigma_1^1$ in Theorem 1.17 is necessary. The issue is that the map $\pi_1(U\widehat{\Sigma}) \to \mathrm{Mod}(\Sigma)$ constructed in Theorem 1.17 is not injective in this case (for instance, this follows from the explicit formulas below). The correct statement is as follows.

THEOREM 1.18. *There is a short exact sequence*

$$1 \longrightarrow \mathbb{Z} \longrightarrow \mathrm{Mod}_1^1 \longrightarrow \mathrm{Mod}_1 \longrightarrow 1,$$

where the kernel $\mathbb{Z}$ is generated by the Dehn twist about the boundary component of Σ_1^1.

For the proof, see [**FM12**, p. 57 & Theorem 3.1.9].

Pushing along simple closed curves. Let Σ and β and $\widehat{\Sigma}$ be as in Theorem 1.17. The kernel $\pi_1(U\widehat{\Sigma}) \subset \mathrm{Mod}(\Sigma)$ of the exact sequence in Theorem 1.17 is known as the *disc-pushing* subgroup. We will occasionally need explicit formulas for elements of it. First, the loop around the fiber of the unit tangent bundle $U\widehat{\Sigma}$ corresponds to the Dehn twist T_β. As far as other elements go, it is easiest to deal with their projections to $\pi_1(\widehat{\Sigma})$. Let $\gamma \in \pi_1(\widehat{\Sigma})$ be an element that can be represented by a simple closed curve. Taking the derivative of a smooth representative of γ, we get a lift $\tilde{\gamma} \in \pi_1(U\widehat{\Sigma})$. Two smooth representatives of γ which are homotopic are smoothly isotopic, so $\tilde{\gamma}$ does not depend on the choice of a smooth representative. We will denote the mapping class in $\mathrm{Mod}(\Sigma)$ associated to $\tilde{\gamma} \in \pi_1(U\widehat{\Sigma}) \subset \mathrm{Mod}(\Sigma)$ by Push_γ. As is shown in Figure 1.7, we can write

$$\mathrm{Push}_\gamma = T_{\tilde{\gamma}_1} T_{\tilde{\gamma}_2}^{-1}$$

for two simple closed curves $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ in Σ that map to the boundary components of a tubular neighborhood of γ in $\widehat{\Sigma}$. The curve $\tilde{\gamma}_1$ lies to the right of γ and the curve $\tilde{\gamma}_2$ lies to the left of γ.

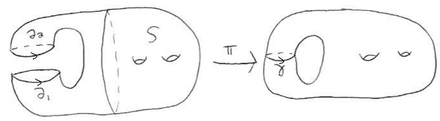

FIGURE 1.8. On the right is a surface Σ_g and an oriented nonseparating simple closed curve γ. On the left is the cut-open surface $\Sigma_{g,\gamma}$ and a γ-splitting surface S

Stabilizer of nonseparating curve. For some $g \geqslant 2$, let γ be an oriented nonseparating simple closed curve on Σ_g. We will use the Birman exact sequence to understand the stabilizer $(\mathrm{Mod}_g)_\gamma$ in Mod_g of γ. Define $\Sigma_{g,\gamma}$ to be the result of cutting Σ_g along γ, let $\mathrm{Mod}_{g,\gamma}$ be the mapping class group of $\Sigma_{g,\gamma}$, and let $\{\partial_1, \partial_2\}$ be the boundary components of $\Sigma_{g,\gamma}$. There is a surjective map $\pi : \mathrm{Mod}_{g,\gamma} \to (\mathrm{Mod}_g)_\gamma$ obtained by gluing ∂_1 and ∂_2 back together (see Figure 1.8)

REMARK 1.19. The map π is surjective because γ is oriented; if it were unoriented, then the image of π would be an index 2 subgroup of $(\mathrm{Mod}_g)_\gamma$.

The map π is not injective; indeed,

$$\pi(T_{\partial_1}) = \pi(T_{\partial_2}) = T_\gamma,$$

so $T_{\partial_1} T_{\partial_1}^{-1} \in \ker(\pi)$. The following lemma says that this is the only thing that goes wrong.

LEMMA 1.20. *For $g \geqslant 2$, let γ be an oriented nonseparating simple closed curve on Σ_g. Let $\{\partial_1, \partial_2\}$ be the boundary components of $\Sigma_{g,\gamma}$. Then there is a short exact sequence*

$$1 \longrightarrow \mathbb{Z} \longrightarrow \mathrm{Mod}_{g,\gamma} \longrightarrow (\mathrm{Mod}_g)_\gamma \longrightarrow 1,$$

where $\mathbb{Z}$ is generated by $T_{\partial_1} T_{\partial_2}^{-1}$.

While the proof of Lemma 1.20 is not hard, it would require a small digression, so we omit it. See [**FM12**, Theorem 3.18] for a proof.

Letting $\widehat{\Sigma}_{g,\gamma}$ be the surface obtained by gluing a disc to $\Sigma_{g,\gamma}$ along ∂_1, Theorem 1.17 says that there is a short exact sequence

$$1 \longrightarrow \pi_1(U\widehat{\Sigma}_{g,\gamma}) \longrightarrow \mathrm{Mod}_{g,\gamma} \longrightarrow \mathrm{Mod}(\widehat{\Sigma}_{g,\gamma}) \longrightarrow 1. \tag{1}$$

In this case, it turns out that the Birman exact sequence splits. A *γ-splitting surface* is a subsurface S of $\Sigma_{g,\gamma}$ such that $\Sigma_{g,\gamma} \setminus \mathrm{Int}(S)$ is a 3-holed sphere two of whose boundary components are ∂_1 and ∂_2 (see Figure 1.8). Letting S be a γ-splitting surface, observe that $S \cong \widehat{\Sigma}_{g,\gamma}$; indeed, regarding S as a subsurface of $\widehat{\Sigma}_{g,\gamma}$ via the inclusion $\Sigma_{g,\gamma} \hookrightarrow \widehat{\Sigma}_{g,\gamma}$, the surface $\widehat{\Sigma}_{g,\gamma}$ deformation retracts onto S. Identifying $\mathrm{Mod}(\widehat{\Sigma}_{g,\gamma})$ with $\mathrm{Mod}(S)$ via this deformation retraction, the map $\mathrm{Mod}(S) \to \mathrm{Mod}_{g,\gamma}$ that extends mapping classes on S by the identity provides a splitting of (1). We summarize this discussion in the following lemma.

LEMMA 1.21. *For $g \geqslant 2$, let γ be an oriented nonseparating simple closed curve on Σ_g. Let $\{\partial_1, \partial_2\}$ be the boundary components of $\Sigma_{g,\gamma}$ and let $\widehat{\Sigma}_{g,\gamma}$ be the surface obtained by gluing a disc to $\Sigma_{g,\gamma}$ along ∂_1. Finally, let S be a γ-splitting surface in $\Sigma_{g,\gamma}$. Then we have a decomposition*

$$\mathrm{Mod}_{g,\gamma} = \pi_1(\widehat{\Sigma}_{g,\gamma}) \rtimes \mathrm{Mod}(S).$$

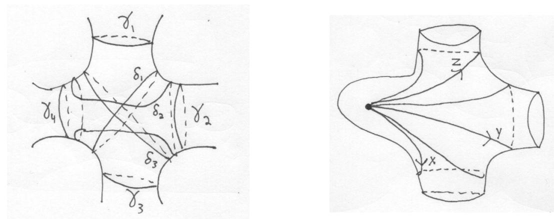

FIGURE 1.9. The left hand figure depicts the curves in involved in the lantern relation $T_{\gamma_1}T_{\gamma_2}T_{\gamma_3}T_{\gamma_4} = T_{\delta_1}T_{\delta_2}T_{\delta_3}$. The right hand figure depicts the relation $xyz = 1$ in $\pi_1(\Sigma_0^3)$ which "lifts" to the lantern relation.

1.5. Some relations between Dehn twists

McCool [**McC75**] proved that the mapping class group is finitely presentable (there is also an influential later proof by Hatcher–Thurston [**HT80**]). See [**FM12**, §5.5.3] for a detailed proof of this. While we will not need to know a complete presentation for the mapping class group, we will need to know three important families of relations.

Disjointness relation. The first says that Dehn twists about disjoint simple closed curves commute.

LEMMA 1.22 (Disjointness relation). *Let Σ be a surface and let γ and γ' be disjoint simple closed curves on Σ. Then T_γ commutes with $T_{\gamma'}$.*

PROOF. Recall that the *support* of a diffeomorphism $F : \Sigma \to \Sigma$ is the closure of the set $\{p \in \Sigma \mid F(p) \neq \text{id}\}$. Our assumptions imply that we can choose diffeomorphisms representing T_γ and $T_{\gamma'}$ whose supports are disjoint. This clearly implies that these diffeomorphisms commute. □

Conjugation. The second reflects the fact that conjugation in $\text{Mod}(\Sigma)$ behaves similarly to conjugation in groups of matrices : it "changes coordinates". We have already used it in the proof of Corollary 1.10.

LEMMA 1.23 (Conjugation relation). *Let Σ be a surface and γ be a simple closed curve on Σ. Then for all $f \in \text{Mod}(\Sigma)$ we have $fT_\gamma f^{-1} = T_{f(\gamma)}$.*

PROOF. Obvious. □

Lantern relation. Our final relation is called the lantern relation. It was first discovered by Dehn [**Deh38**] in 1938, but was forgotten until it was rediscovered by Johnson [**Joh79**] in 1979. It is the most important relation for the study of the Torelli group, and we will use it many times. The proof we will give was discovered independently by Margalit–McCammond [**MM09**] and Putman [**Put09**].

LEMMA 1.24 (Lantern relation). *Let Σ be a surface and let γ_1, γ_2, γ_3, γ_4, δ_1, δ_2, and δ_3 be simple closed curves on Σ which are arranged like the curves in Figure 1.9. Then*

$$T_{\gamma_1}T_{\gamma_2}T_{\gamma_3}T_{\gamma_4} = T_{\delta_1}T_{\delta_2}T_{\delta_3}.$$

PROOF. The purported relation is supported on a 4-holed sphere, so we can assume without loss of generality that $\Sigma = \Sigma_0^4$ and that $\gamma_1, \ldots, \gamma_4$ are the boundary components of Σ. Let $\widehat{\Sigma}$ be the result of gluing a disc to γ_4 and let $x, y, z \in \pi_1(\widehat{\Sigma})$ be the curves depicted in Figure 1.9, so $xyz = 1$. The relation $xyz = 1$ in $\pi_1(\widehat{\Sigma})$ lifts to a relation

$$\text{Push}_z \text{Push}_y \text{Push}_x = T_\beta^k \tag{2}$$

in $\pi_1(U\widehat{\Sigma}) \subset \text{Mod}(\Sigma)$ for some $k \in \mathbb{Z}$; the T_β^k appears because the loop around the fiber generates the kernel of the projection $\pi_1(U\widehat{\Sigma}) \to \pi_1(\widehat{\Sigma})$, and the order is reversed because functions are composed right-to-left while curves in the fundamental group are composed left-to-right. Observe that

$$\text{Push}_x = T_{\delta_3} T_{\gamma_3}^{-1} \quad \text{and} \quad \text{Push}_y = T_{\delta_2} T_{\gamma_2}^{-1} \quad \text{and} \quad \text{Push}_z = T_{\delta_1} T_{\gamma_1}^{-1}.$$

Lemma 1.22 implies that the T_{γ_i} commute with each other and with the T_{δ_j}. We can thus rearrange the terms in (2) to get

$$T_{\gamma_1} T_{\gamma_2} T_{\gamma_3} T_{\gamma_4}^k = T_{\delta_1} T_{\delta_2} T_{\delta_3}.$$

It is an easy exercise to see that $k = 1$. $\square$

CHAPTER 2

The symplectic representation

The Torelli group is the kernel of the action of Mod_g on $\mathrm{H}_1(\Sigma_g;\mathbb{Z})$. This chapter is devoted to a preliminary study of this action. As we will see, it induces a surjective representation from Mod_g to $\mathrm{Sp}_{2g}(\mathbb{Z})$. As notation, if γ is an oriented simple closed curve on a surface Σ, then $[\gamma]$ will denote the associated element of $\mathrm{H}_1(\Sigma;\mathbb{Z})$.

2.1. The algebraic intersection form

We already met the algebraic intersection form $\hat{i}(\cdot,\cdot)$ when we proved that $\mathrm{Mod}(\mathbb{T}^2) \cong \mathrm{SL}_2(\mathbb{Z})$ (Proposition 1.6).

Nondegeneracy. For a closed surface Σ_g, Poincaré duality implies that $\hat{i}(\cdot,\cdot)$ is *nondegenerate* in the sense that the map

$$\mathrm{H}_1(\Sigma_g;\mathbb{Z}) \longrightarrow (\mathrm{H}_1(\Sigma_g;\mathbb{Z}))^*$$
$$h \mapsto (x \mapsto \hat{i}(h,x))$$

is an isomorphism. Here $(\mathrm{H}_1(\Sigma_g;\mathbb{Z}))^*$ is the dual $\mathbb{Z}$-module $\mathrm{Hom}(\mathrm{H}_1(\Sigma_g;\mathbb{Z}),\mathbb{Z})$; of course, $(\mathrm{H}_1(\Sigma_g;\mathbb{Z}))^* \cong \mathrm{H}^1(\Sigma_g;\mathbb{Z})$. Since the map $\mathrm{H}_1(\Sigma_g^1;\mathbb{Z}) \to \mathrm{H}_1(\Sigma_g;\mathbb{Z})$ is an isomorphism, it follows that $\hat{i}(\cdot,\cdot)$ is also nondegenerate for a surface with one boundary component.

Summary. This is summarized in the following lemma. A *symplectic form* on a free finite-rank $\mathbb{Z}$-module M is a nondegenerate alternating bilinear form on M.

LEMMA 2.1. *For $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, the algebraic intersection form $\hat{i}(\cdot,\cdot)$ on $\mathrm{H}_1(\Sigma_g^b;\mathbb{Z})$ is a symplectic form.*

REMARK 2.2. Lemma 2.1 is false for $b \geqslant 2$; indeed, if $b \geqslant 2$ and β is an oriented boundary component of Σ_g^b, then $[\beta] \neq 0$ but $\hat{i}([\beta],h) = 0$ for all $h \in \mathrm{H}_1(\Sigma_g^b;\mathbb{Z})$. In other words, $[\beta]$ is a nonzero element of the kernel of the map $\mathrm{H}_1(\Sigma_g^b;\mathbb{Z}) \to (\mathrm{H}_1(\Sigma_g^b;\mathbb{Z}))^*$ discussed above, so $\hat{i}(\cdot,\cdot)$ is degenerate.

Symplectic basis. If M is a free finite-rank $\mathbb{Z}$-module equipped with a symplectic form $\omega(\cdot,\cdot)$, then a *symplectic basis* for M is a free basis $\{a_1, b_1, \ldots, a_g, b_g\}$ for M such that

$$\omega(a_i,a_j) = \omega(b_i,b_j) = 0 \quad \text{and} \quad \omega(a_i,b_j) = \delta_{ij}$$

for all $1 \leqslant i,j \leqslant g$. For example, if $\alpha_1, \beta_1, \ldots, \alpha_g, \beta_g$ are the oriented simple closed curves on Σ_g^1 depicted in Figure 2.1, then $\{[\alpha_1],[\beta_1],\ldots,[\alpha_g],[\beta_g]\}$ is a symplectic basis for $\mathrm{H}_1(\Sigma_g^1;\mathbb{Z})$. The following lemma says that these always exist; it implies in particular that the rank of M is always even.

LEMMA 2.3. *Let M be a free finite-rank $\mathbb{Z}$-module equipped with a symplectic form $\omega(\cdot,\cdot)$. Then M has a symplectic basis.*

PROOF. The proof will be by induction on the rank n of M. The base case is $n = 0$, where the lemma is trivial. Now assume that $n > 0$ and that the lemma is true for all smaller n. Let $\{x_1, \ldots, x_n\}$ be a free basis for M. The nondegeneracy of $\omega(\cdot,\cdot)$ implies that there exists some $b_1 \in M$ such that

$$\omega(x_i, b_1) = \delta_{1i} \qquad (1 \leqslant i \leqslant n).$$

Set $a_1 = x_1$, so $\omega(a_1, b_1) = 1$. Define a surjective homomorphism $\pi : M \to \langle a_1, b_1\rangle$ via the formula

$$\pi(x) = \omega(x, b_1) \cdot a_1 - \omega(x, a_1) \cdot b_1.$$

Clearly π is a split surjection whose kernel is the *orthogonal complement* of $\langle a_1, b_1\rangle$, that is, the set

$$\langle a_1, b_1\rangle^{\perp} := \{x \in M \mid \omega(a_1, x) = \omega(b_1, x) = 0\}.$$

It follows that $M = \langle a_1, b_1\rangle \oplus \langle a_1, b_1\rangle^{\perp}$. It is easy to see that the restriction of ω to $\langle a_1, b_1\rangle^{\perp}$ is still nondegenerate, so by induction $\langle a_1, b_1\rangle^{\perp}$ has a symplectic basis $\{a_2, b_2, \ldots, a_g, b_g\}$ for some $g \geqslant 0$. The desired symplectic basis for M is then $\{a_1, b_1, \ldots, a_g, b_g\}$. □

COROLLARY 2.4 (Uniqueness of symplectic forms). *Let M and M' be a rank $2g$ free $\mathbb{Z}$-modules equipped with symplectic forms $\omega(\cdot,\cdot)$ and $\omega'(\cdot,\cdot)$, respectively. Then there exists a $\mathbb{Z}$-linear isomorphism $\psi : M \to M'$ such that*

$$\omega(x, y) = \omega'(\psi(x), \psi(y)) \qquad (x, y \in M).$$

PROOF. Let $\{a_1, b_1, \ldots, a_g, b_g\}$ (resp. $\{a'_1, b'_1, \ldots, a'_g, b'_g\}$) be a symplectic basis for M (resp. M') with respect to $\omega(\cdot,\cdot)$ (resp. $\omega'(\cdot,\cdot)$). The isomorphism ψ is then defined via the formulas

$$\psi(a_i) = a'_i \quad \text{and} \quad \psi(b_i) = b'_i$$

for $1 \leqslant i \leqslant g$. □

2.2. The symplectic representation : statement of surjectivity

We now introduce the symplectic representation of the mapping class group.

The symplectic group. The *genus g symplectic group*, denoted $\mathrm{Sp}_{2g}(\mathbb{Z})$, is defined as follows. Let M be a rank $2g$ free $\mathbb{Z}$-module equipped with a symplectic form $\omega(\cdot,\cdot)$; for instance, M might be $\mathrm{H}_1(\Sigma_g;\mathbb{Z})$ and $\omega(\cdot,\cdot)$ might be the algebraic intersection pairing. Define

$$\mathrm{Sp}(M,\omega) = \{\phi \in \mathrm{GL}(M) \mid \omega(x, y) = \omega(\phi(x), \phi(y)) \text{ for all } x, y \in M\} \subset \mathrm{GL}(M)$$

Then $\mathrm{Sp}_{2g}(\mathbb{Z})$ is the subgroup of $\mathrm{GL}_{2g}(\mathbb{Z})$ obtained by considering $\mathrm{Sp}(M,\omega)$ with respect to a symplectic basis for M. Corollary 2.4 implies that $\mathrm{Sp}_{2g}(\mathbb{Z})$ is well-defined.

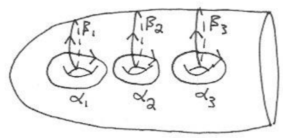

FIGURE 2.1. A geometric symplectic basis on Σ_3^1

REMARK 2.5. It is also common to order the symplectic basis as

$$\{a_1, a_2, \dots, a_g, b_1, b_2, \dots, b_g\}$$

when defining $\mathrm{Sp}_{2g}(\mathbb{Z})$.

REMARK 2.6. The calculation we performed right before Proposition 1.6 can be interpreted as saying that $\mathrm{Sp}_2(\mathbb{Z}) = \mathrm{SL}_2(\mathbb{Z})$.

The mapping class group. If Σ is a surface, then the action of $\mathrm{Mod}(\Sigma)$ on $\mathrm{H}_1(\Sigma;\mathbb{Z})$ preserves the algebraic intersection form. For $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, Lemma 2.1 implies that this yields a representation

$$\mathrm{Mod}_g^b \longrightarrow \mathrm{Sp}_{2g}(\mathbb{Z}).$$

The following theorem says that this is surjective; it is the main result of this chapter. It was originally proved by Burkhardt in 1890 [**Bur90**, pp. 209–212], who wrote down mapping classes that map to generators of $\mathrm{Sp}_{2g}(\mathbb{Z})$ that were previously found by Clebsch–Gordan [**CG66**]. The first modern proof is due to Meeks-Patrusky [**MP78**, Theorem 2], and our proof is a variant of theirs.

THEOREM 2.7. *For $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, the map $\mathrm{Mod}_g^b \to \mathrm{Sp}_{2g}(\mathbb{Z})$ coming from the action of Mod_g^b on $\mathrm{H}_1(\Sigma_g^b;\mathbb{Z})$ is surjective.*

The proof of Theorem 2.7 is in §2.5. This is prefaced with two sections of preliminaries.

2.3. Realizing primitive homology classes

The first step towards proving Theorem 2.7 is to determine which elements of $\mathrm{H}_1(\Sigma_g^b;\mathbb{Z})$ can be realized by simple closed curves. An element a of a free abelian group A is *primitive* if we cannot write $a = n \cdot a'$ for some $a' \in A$ and $n \in \mathbb{Z}$ satisfying $n \geqslant 2$. The earliest proofs of the following theorem that we are aware of are due to Schafer [**Sch76**] (who actually deduced it from Theorem 2.7!) and Meyerson [**Mey76**]. Our proof is inspired by Meeks–Patrusky [**MP78**].

THEOREM 2.8. *For $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, consider some nonzero $h \in \mathrm{H}_1(\Sigma_g^b;\mathbb{Z})$. Then there exists an oriented simple closed curve γ on Σ_g^b with $[\gamma] = h$ if and only if h is primitive.*

REMARK 2.9. If $b \geqslant 2$, then not all primitive elements of $\mathrm{H}_1(\Sigma_g^b;\mathbb{Z})$ can be realized by oriented simple closed curves. See [**MP78**] for how to correct the statement in this case.

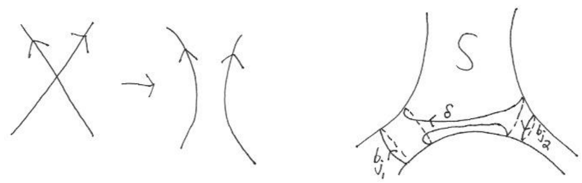

FIGURE 2.2. The left hand figure shows how to resolve the intersections and self-intersection of the γ_i'; the result is a collection of disjoint oriented simple closed curves. In the right hand figure, the surface S lies on the same size of b_{j_1} and b_{j_2}. The simple closed curves δ on S satisfies $[\delta] = [b_{j_1}] + [b_{j_2}]$.

PROOF OF THEOREM 2.8. First assume that such a γ exists and that $h = n \cdot h'$ for some $h' \in \mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$. We will prove that $n = \pm 1$. Separating curves are nullhomologous, so γ does not separate Σ_g^b (this uses the fact that $0 \leqslant b \leqslant 1$). This implies that there exists an oriented simple closed curve δ on Σ_g^b that intersects γ once. Observe that $\hat{i}([\gamma], [\delta]) = \pm 1$. We conclude that

$$\pm 1 = \hat{i}([\gamma], [\delta]) = \hat{i}(n \cdot h', [\delta]) = n \cdot \hat{i}(h', [\delta]),$$

so $n = \pm 1$. We remark that this argument first appeared in [**Mey76**], which attributes it to Samelson.

Now assume that h is primitive. We will construct γ in two steps.

STEP 1. *There exist disjoint oriented simple closed curves $\gamma_1, \ldots, \gamma_k$ in the interior of Σ_g^b such that $h = [\gamma_1] + \cdots + [\gamma_k]$.*

Clearly we can write $h = [\gamma_1'] + \cdots + [\gamma_\ell']$ for some oriented curves $\gamma_1', \ldots, \gamma_\ell'$ in the interior of Σ_g^b (not necessarily simple or disjoint). As is shown in Figure 2.2, we can then "resolve" the intersections and self-intersections of the γ_i' to obtain the desired set of disjoint oriented simple closed curves.

STEP 2. *There exists an oriented simple closed curve γ such that $[\gamma] = h$.*

Using Step 1, we can write

$$h = [\gamma_1] + \cdots + [\gamma_k] \qquad (3)$$

for some disjoint oriented simple closed curves $\gamma_1, \ldots, \gamma_k$ in the interior of Σ_g^b. Choose (3) such that k is as small as possible. We will prove that $k = 1$.

Consider any component S of the result of cutting Σ_g^b along the γ_i. Let $b_1, \ldots, b_m$ be the boundary components of S that lie in the interior of Σ_g^b, so each b_i equals γ_{j_i} for some $1 \leqslant j_i \leqslant k$ (the j_i are not necessarily distinct); orient b_i using the orientation of γ_{j_i}. If $m = 1$, then b_1 is nullhomologous (this uses the fact that $0 \leqslant b \leqslant 1$), so we can discard γ_{i_1} from (3), contradicting the minimality of k. Thus $m > 1$. If $1 \leqslant j_1 < j_2 \leqslant m$ are such that the interior of S lies on the same side of b_{j_1} and b_{j_2}, then $\gamma_{i_{j_1}} \neq \gamma_{i_{j_2}}$ and as shown in Figure 2.2 we can replace $\gamma_{i_{j_1}}$ and $\gamma_{i_{j_2}}$ in (3) with a single oriented curve δ, again contradicting the minimality of k. We conclude that the interior of S must lie on *different* sides of b_{j_1} and b_{j_2} for all $1 \leqslant j_1 < j_2 \leqslant m$.

The upshot of the above considerations is that $m = 2$ and that S lies on different sides of b_1 and b_2. This implies that $[\gamma_{i_1}] = [\gamma_{i_2}]$ (here we are again using the fact that $0 \leqslant b \leqslant 1$). Since this holds for *all* components S of the result of cutting Σ_g^b along the γ_i, we deduce that $[\gamma_i] = [\gamma_j]$ for all $1 \leqslant i, j \leqslant k$, so $h = k \cdot [\gamma_1]$. Since h is primitive, we conclude that $k = 1$. □

2.4. Realizing symplectic bases

We now extend Theorem 2.8 to certain systems of curves.

Geometric symplectic bases. Fix $g \geqslant 0$ and $0 \leqslant b \leqslant 1$. A *geometric symplectic basis* on Σ_g^b is a set $\{\alpha_1, \beta_1, \ldots, \alpha_g, \beta_g\}$ of oriented simple closed curves on Σ_g^b with the following two properties.

- For $1 \leqslant i \leqslant g$, the curves α_i and β_i intersect once with a positive sign.
- All other pairs of distinct curves in the set are disjoint.

See Figure 2.1. Observe that if $\{\alpha_1, \beta_1, \ldots, \alpha_g, \beta_g\}$ is a geometric symplectic basis on Σ_g^b, then $\{[\alpha_1], [\beta_1], \ldots, [\alpha_g], [\beta_g]\}$ is a symplectic basis for $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$.

Realizing symplectic bases. The following proposition says that any symplectic basis can be realized by a geometric symplectic basis.

Proposition 2.10. *For $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, consider a symplectic basis $\{a_1, b_1, \ldots, a_g, b_g\}$ for $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$. There is then a geometric symplectic basis $\{\alpha_1, \beta_1, \ldots, \alpha_g, \beta_g\}$ on Σ_g^b satisfying $[\alpha_i] = a_i$ and $[\beta_i] = b_i$ for $1 \leqslant i \leqslant g$.*

For the proof of Proposition 2.10, we will need the following lemma.

Lemma 2.11. *Let Σ be a surface and let α be an oriented nonseparating simple closed curve on Σ. Set $a = [\alpha]$ and let $b \in \mathrm{H}_1(\Sigma; \mathbb{Z})$ satisfy $\hat{i}(a, b) = 1$. Then we can find an oriented simple closed curve β on Σ that intersects α once and satisfies $[\beta] = b$.*

Proof. We divide the proof into two steps.

Step 1. *We can find an oriented closed curve β' on Σ (not necessarily simple) that intersects α once and satisfies $[\beta'] = b$.*

Since α is nonseparating, we can find an oriented simple closed curve ν' on Σ that intersects α once with $\hat{i}(a, [\nu']) = 1$. Let $T \subset \Sigma$ be a closed regular neighborhood of $\alpha \cup \nu'$, so T is homeomorphic to a 1-holed torus. Also, let $S = \Sigma \setminus \mathrm{Int}(T)$. We then have $\mathrm{H}_1(\Sigma; \mathbb{Z}) = U \oplus V$, where U and V are the images of $\mathrm{H}_1(T; \mathbb{Z})$ and $\mathrm{H}_1(S; \mathbb{Z})$ in $\mathrm{H}_1(\Sigma; \mathbb{Z})$, respectively. Write $b = u + v$ with $u \in U$ and $v \in V$. We thus have $\hat{i}(a, u) = 1$ and $\hat{i}(a, v) = 0$.

The set $\{a, [\nu']\}$ is a basis for U; write $u = p \cdot a + q \cdot [\nu']$. Since

$$1 = \hat{i}(a, u) = \hat{i}(a, p \cdot a + q \cdot [\nu']) = q,$$

we have $q = 1$. Setting $\nu = T_\alpha^p(\nu')$, the oriented simple closed curve ν lies in T, intersects α once, and satisfies $[\nu] = u$.

Let η be an oriented closed curve on S such that $[\eta] = v$. The homology class v need not be primitive, so we might not be able to choose η to be simple.

We have $b = [\nu] + [\eta]$. The desired closed curve β' can be obtained by "band-summing" the curves ν and η as shown in Figure 2.3.

Step 2. *We can find an oriented simple closed curve β on Σ that intersects α once and satisfies $[\beta] = b$.*

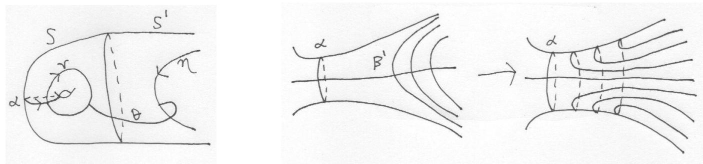

FIGURE 2.3. The left hand figure shows the entire simple closed curve ν and a segment of the (not necessarily simple) closed curve η. Let θ be an arc which starts on the left side of ν and ends on the right side of η and is disjoint from α (see the figure; as shown here, it might not be possible for θ to be disjoint from η). We can then "band-sum" ν and η by removing small segments of ν and η around the endpoints of θ and gluing in two parallel copies of θ; the result is a (not necessarily simple) closed curve β' satisfying $[\beta'] = [\nu] + [\eta]$. The right hand figure shows the segment of β' that intersects α plus some segments of β' that intersect β'. As is shown there, we can "comb" those intersections towards the intersection point of α and β' and then "pass them through α"; the result is a simple closed curve β'' that intersects α once.

As shown in Figure 2.3.b, we can "comb" all of the self-intersections of β' to the point $\beta' \cap \alpha$ and then "pass them through α". The result is a simple closed curve β'' on Σ that intersects α once. However, we no longer have $[\beta''] = b$; instead, we have $[\beta''] = b + n \cdot a$ for some $n \in \mathbb{Z}$. The desired simple closed curve β is then $T_\alpha^{-n}(\beta'')$. □

PROOF OF PROPOSITION 2.10. The proof will be by induction on g. The base case is $g = 0$, where the theorem is vacuous. Assume now that $g > 0$. Using Theorem 2.8 and Lemma 2.11, we can find oriented simple closed curves α_1 and β_1 on Σ_g^b that intersect once and satisfy $[\alpha_1] = a_1$ and $[\beta_1] = b_1$. Let T be a closed regular neighborhood of $\alpha_1 \cup \beta_1$, so T is a one-holed torus. Define $S = \Sigma_g^b \setminus \mathrm{Int}(T)$. We then have a decomposition $\mathrm{H}_1(\Sigma_g^b;\mathbb{Z}) = U \oplus V$, where U is the image of $\mathrm{H}_1(T;\mathbb{Z})$ in $\mathrm{H}_1(\Sigma_g^b;\mathbb{Z})$ and V is the image of $\mathrm{H}_1(S;\mathbb{Z})$. This decomposition is orthogonal with respect to $\hat{i}(\cdot,\cdot)$. Also, $a_1, b_1 \in U$. It follows that $a_2, b_2, \ldots, a_g, b_g \in V$.

The kernel of the map $\mathrm{H}_1(S;\mathbb{Z}) \to \mathrm{H}_1(\Sigma_g^b;\mathbb{Z})$ (possibly 0) is generated by the homology class of $\partial T \subset S$. Let $\widehat{S}$ be the result of gluing a disc to S along ∂T. There is then a natural isomorphism $V \cong \mathrm{H}_1(\widehat{S};\mathbb{Z})$. Using this identification, we can identify $a_2, b_2, \ldots, a_g, b_g$ with elements of $\mathrm{H}_1(\widehat{S};\mathbb{Z})$. The surface $\widehat{S}$ has the same number of boundary components as Σ_g^b (i.e. at most 1), so by induction we can find a geometric symplectic basis $\{\widehat{\alpha}_2, \widehat{\beta}_2, \ldots, \widehat{\alpha}_g, \widehat{\beta}_g\}$ on $\widehat{S}$ satisfying $[\widehat{\alpha}_i] = a_i$ and $[\widehat{\beta}_i] = b_i$ for $2 \leqslant i \leqslant g$. Homotoping the $\widehat{\alpha}_i$ and $\widehat{\beta}_i$, we can assume that they are disjoint from the disc that was glued to S to form $\widehat{S}$. They thus are the images of curves $\alpha_2, \beta_2, \ldots, \alpha_g, \beta_g$ in $S \subset \Sigma_g^b$ under the inclusion $S \hookrightarrow \widehat{S}$. Clearly $\{\alpha_1, \beta_1, \ldots, \alpha_g, \beta_g\}$ is the desired geometric symplectic basis. □

2.5. Proof of surjectivity

We are finally ready to prove Theorem 2.7.

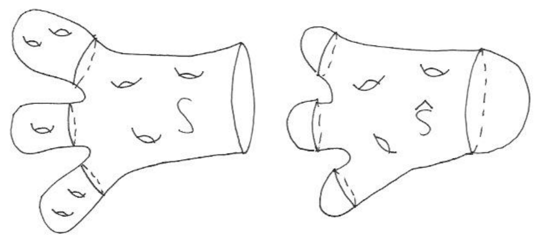

FIGURE 2.4. All the boundary components of the subsurface S of Σ_g^b are separating curves; the image of $\mathrm{H}_1(S;\mathbb{Z})$ in $\mathrm{H}_1(\Sigma_g^b;\mathbb{Z})$ is isomorphic to $\mathrm{H}_1(\widehat{S};\mathbb{Z})$, where $\widehat{S}$ is the result of gluing discs to S along all of its boundary components.

PROOF OF THEOREM 2.7. Let $\{\alpha_1, \beta_1, \ldots, \alpha_g, \beta_g\}$ be a geometric symplectic basis on Σ_g^b. Set $a_i = [\alpha_i]$ and $b_i = [\beta_i]$ for $1 \leqslant i \leqslant g$, so $\{a_1, b_1, \ldots, a_g, b_g\}$ is a symplectic basis for $\mathrm{H}_1(\Sigma_g^b;\mathbb{Z})$. Consider $M \in \mathrm{Sp}_{2g}(\mathbb{Z})$. Observe that

$$M(a_1), M(b_1), \ldots, M(a_g), M(b_g)$$

is a symplectic basis for $\mathrm{H}_1(\Sigma_g^b;\mathbb{Z})$, so Proposition 2.10 says that there exists a geometric symplectic basis $\{\alpha_1', \beta_1', \ldots, \alpha_g', \beta_g'\}$ on Σ_g^b such that $[\alpha_i'] = M(a_i)$ and $[\beta_i'] = M(b_i)$ for $1 \leqslant i \leqslant g$. Using the classification of surface trick (see §1.3), we can find some $f \in \mathrm{Mod}_g^b$ such that $f(\alpha_i) = \alpha_i'$ and $f(\beta_i) = \beta_i'$ for $1 \leqslant i \leqslant g$. Clearly $\psi(f) = M$. □

2.6. Variants on realizing symplectic bases

This section is a bit of detour : it contains some variants on the results in §2.4 that will be needed in subsequent chapters. The reader is advised to skip reading it the first time they read this chapter.

Realizing a symplectic basis on a subsurface. If M is a free $\mathbb{Z}$-module equipped with a symplectic form and $V \subset M$ is a $\mathbb{Z}$-submodule, then we will say that V is a *symplectic subspace* of M if the restriction of the symplectic form on M to V is symplectic. Fix $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, and let $S \subset \Sigma_g^b$ be a subsurface. If S has multiple boundary components, then its algebraic intersection form is degenerate and the image of $\mathrm{H}_1(S;\mathbb{Z})$ in $\mathrm{H}_1(\Sigma_g^b;\mathbb{Z})$ need not be symplectic. However, if all of the boundary components of S are separating (i.e. nullhomologous) curves in Σ_g^b, then the image of $\mathrm{H}_1(S;\mathbb{Z})$ in $\mathrm{H}_1(\Sigma_g^b;\mathbb{Z})$ is symplectic. Indeed, letting $\widehat{S}$ be the closed surface obtained by gluing discs to all of the boundary components of S, the image of $\mathrm{H}_1(S;\mathbb{Z})$ in $\mathrm{H}_1(\Sigma_g^b;\mathbb{Z})$ is easily seen to be isomorphic to $\mathrm{H}_1(\widehat{S};\mathbb{Z})$; see Figure 2.4. This leads us to the following result.

PROPOSITION 2.12. *For $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, let $S \subset \Sigma_g^b$ be a subsurface. Assume that all of the boundary components of S are separating curves in Σ_g^b and let $U \subset \mathrm{H}_1(\Sigma_g^b;\mathbb{Z})$ be the image of $\mathrm{H}_1(S;\mathbb{Z})$ in $\mathrm{H}_1(\Sigma_g^b;\mathbb{Z})$. Finally, let $\{a_1, b_1, \ldots, a_h, b_h\}$ be a symplectic basis for U. Then there exist oriented simple closed curves $\{\alpha_1, \beta_1, \ldots, \alpha_h, \beta_h\}$ in $S \subset \Sigma_g^b$ with the following properties.*

- *$[\alpha_i] = a_i$ and $[\beta_i] = b_i$ for all $1 \leqslant i \leqslant h$, where these homology classes are in $\mathrm{H}_1(\Sigma_g^b;\mathbb{Z})$.*

- *For $1 \leqslant i \leqslant h$, the curves α_i and β_i intersect once.*
- *All other pairs of distinct curves in the set are disjoint.*

PROOF. Let $\widehat{S}$ be the closed genus h surface obtained by gluing discs to all of the boundary components of S, so $U \cong \mathrm{H}_1(\widehat{S};\mathbb{Z})$. Proposition 2.10 says that there exists a geometric symplectic basis $\{\alpha_1, \beta_1, \ldots, \alpha_h, \beta_h\}$ on $\widehat{S}$ such that $[\alpha_i] = a_i$ and $[\beta_i] = b_i$ for $1 \leqslant i \leqslant h$. Homotoping the α_i and β_i, we can assume that they are disjoint from all the glued-on discs, and thus lie in S. □

Completing a partial geometric symplectic basis. For $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, a *partial geometric symplectic basis* on Σ_g^b is a set $\{\alpha_1, \ldots, \alpha_h, \beta_1, \ldots, \beta_{h'}\}$ of oriented simple closed curves on Σ_g^b with the following properties.

- For $1 \leqslant i \leqslant \min(h, h')$, the curves α_i and β_i intersect once with a positive sign.
- All other pairs of distinct curves in the set are disjoint.

We then have the following generalization of Proposition 2.10.

PROPOSITION 2.13. *For $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, consider a symplectic basis $\{a_1, b_1, \ldots, a_g, b_g\}$ for $\mathrm{H}_1(\Sigma_g^b;\mathbb{Z})$. Assume that $\{\alpha_1, \ldots, \alpha_h, \beta_1, \ldots, \beta_{h'}\}$ is a partial geometric symplectic basis on Σ_g^b with $[\alpha_i] = a_i$ for $1 \leqslant i \leqslant h$ and $[\beta_j] = b_j$ for $1 \leqslant j \leqslant h'$. We can then extend our partial geometric symplectic basis to a geometric symplectic basis $\{\alpha_1, \beta_1, \ldots, \alpha_g, \beta_g\}$ on Σ_g^b satisfying $[\alpha_i] = a_i$ and $[\beta_i] = b_i$ for $1 \leqslant i \leqslant g$.*

For the proof of Proposition 2.13, we will need the following generalization of Lemma 2.11.

LEMMA 2.14. *Let Σ be a surface and let $\{\alpha_1, \ldots, \alpha_h\}$ be disjoint oriented simple closed curves on Σ such that $\alpha_1 \cup \cdots \cup \alpha_h$ does not separate Σ. For $1 \leqslant i \leqslant h$, set $a_i = [\alpha_i]$. Let $b_1 \in \mathrm{H}_1(\Sigma;\mathbb{Z})$ satisfy $\hat{i}(a_1, b_1) = 1$ and $\hat{i}(a_i, b_1) = 0$ for $2 \leqslant i \leqslant h$. Then we can find an oriented simple closed curve β_1 on Σ that intersects α_1 once, is disjoint from α_i for $2 \leqslant i \leqslant h$, and satisfies $[\beta_1] = b_1$.*

PROOF. Let $\Sigma' \subset \Sigma$ be the complement of an open regular neighborhood of $\alpha_2 \cup \cdots \cup \alpha_h$; if $\Sigma \cong \Sigma_g^b$, then $\Sigma' \cong \Sigma_{g-h}^{b+2h}$. Letting α_1' be α_1 regarded as a curve on Σ', the curve α_1' does not separate Σ'. Moreover, it is an easy exercise to see that the image of $\mathrm{H}_1(\Sigma';\mathbb{Z})$ in $\mathrm{H}_1(\Sigma;\mathbb{Z})$ is

$$\langle a_2, \ldots, a_h \rangle^{\perp} = \{h \in \mathrm{H}_1(\Sigma;\mathbb{Z}) \mid \hat{i}(a_i, h) = 0 \text{ for } 2 \leqslant i \leqslant h\}.$$

In particular, we can find some $b_1' \in \mathrm{H}_1(\Sigma';\mathbb{Z})$ that maps to $b_1 \in \mathrm{H}_1(\Sigma;\mathbb{Z})$. We have $\hat{i}([\alpha_1'], b_1') = 1$, so Lemma 2.11 implies that we can find an oriented simple closed curve β_1' in Σ' that intersects α_1' once and satisfies $[\beta_1'] = b_1$. Letting β_1 be the image of β_1' in Σ under the inclusion $\Sigma' \hookrightarrow \Sigma$, the curve β_1 satisfies the desired properties. □

PROOF OF PROPOSITION 2.13. The proof is a straightforward generalization of the proof of Proposition 2.10, with Lemma 2.14 used in place of 2.11. □

A variant on finding β. The final result we need is the following variant of Lemma 2.11. In it, the pair of curves $\{\alpha, \alpha'\}$ form what we will call a *bounding pair* in Chapter 3.

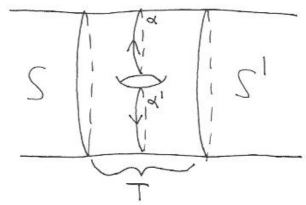

FIGURE 2.5. A bounding pair $\{\alpha, \alpha'\}$ contained in a 2-holed torus T.

LEMMA 2.15. *Let Σ be a surface and let α and α' be disjoint oriented simple closed nonseparating curves such that $[\alpha] = [\alpha']$. Set $a = [\alpha] = [\alpha']$, and let $b \in \mathrm{H}_1(\Sigma; \mathbb{Z})$ satisfy $\hat{i}(a, b) = 1$. Then we can find an oriented simple closed curve β on Σ that intersects both α and α' once and satisfies $[\beta] = b$.*

PROOF. The proof is very similar to the proof of Lemma 2.11. There are two steps.

STEP 1. *We can find an oriented closed curve β' on Σ (not necessarily simple) that intersects both α and α' once and satisfies $[\beta'] = b$.*

Since $[\alpha] = [\alpha']$, the union $\alpha \cup \alpha'$ must separate the surface. As is shown in Figure 2.5, we can then find a 2-holed torus $T \subset \Sigma$ containing $\alpha \cup \alpha'$. Let S be the component of $\Sigma \setminus \mathrm{Int}(T)$ lying to the left of α and S' be the component lying to the right. We then have $\mathrm{H}_1(\Sigma; \mathbb{Z}) = U \oplus V \oplus V'$, where U and V and V' are the images of $\mathrm{H}_1(T; \mathbb{Z})$ and $\mathrm{H}_1(S; \mathbb{Z})$ and $\mathrm{H}_1(S'; \mathbb{Z})$ in $\mathrm{H}_1(\Sigma; \mathbb{Z})$, respectively. Write $b = u + v + v'$ with $u \in U$ and $v \in V$ and $v' \in V'$. We thus have $\hat{i}(a, u) = 1$ and $\hat{i}(a, v) = \hat{i}(a, v') = 0$.

Let ν' be an oriented simple closed curve on T that intersects both α and α' once with a positive sign. The set $\{a, [\nu']\}$ is a basis for U; write $u = p \cdot a + q \cdot [\nu']$. Since

$$1 = \hat{i}(a, u) = \hat{i}(a, p \cdot a + q \cdot [\nu']) = q,$$

we have $q = 1$. Setting $\nu = T_\alpha^p(\nu')$, the oriented simple closed curve ν lies in T, intersects both α and α' once, and satisfies $[\nu] = u$.

Let η be an oriented closed curve on S such that $[\eta] = v$. The homology class v need not be primitive, so we might not be able to choose η to be simple. Similarly, let η' be an oriented closed curve on S' such that $[\eta'] = v'$.

We have $b = [\nu] + [\eta] + [\eta']$. Just like in the proof of Lemma 2.11, the desired closed curve β' can be obtained by "band-summing" the curves ν and η and η'.

STEP 2. *We can find an oriented simple closed curve β on Σ that intersects both α and α' once and satisfies $[\beta] = b$.*

Like like we did in the proof of Lemma 2.11, we can "comb" the self-intersections of β' to α and them "pass them through α" (some intersections will be on the left of α and some on the right; we avoid combing intersections through α'). The result is an oriented simple closed curve β'' on Σ that intersects both α and α' once and satisfies $[\beta''] = b + n \cdot a$ for some $n \in \mathbb{Z}$. The desired oriented simple closed curve β is then $T_\alpha^{-n}(\beta'')$. □

CHAPTER 3

Basic properties of the Torelli group

We finally come to the Torelli group. This chapter defines the Torelli group and discusses a number of its basic properties.

3.1. Definition and low-complexity examples

We start by defining the Torelli group.

Definition. For $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, the *Torelli group*, denoted $\mathcal{I}_g^b$, is the kernel of the action of Mod_g^b on $\text{H}_1(\Sigma_g^b; \mathbb{Z})$. Using Theorem 2.7, we have an exact sequence

$$1 \longrightarrow \mathcal{I}_g^b \longrightarrow \text{Mod}_g^b \longrightarrow \text{Sp}_{2g}(\mathbb{Z}) \longrightarrow 1.$$

Just like for Mod_g^b, we will often omit the b when it vanishes. We will also use the notation $\mathcal{I}(\Sigma)$ for the Torelli group of a surface Σ. We emphasize that Σ is allowed to have at most 1 boundary component.

REMARK 3.1. The reader probably wonders why we have not defined the Torelli group on a surface with multiple boundary components. The issue is that it is not clear what the correct definition should be. This is related to the fact (see Remark 2.2) that on these surfaces the algebraic intersection pairing is degenerate. One obvious thing to do would be to simply define it to be the kernel of the action of the mapping class group on homology, but this turns out not to be particularly well-behaved. See [**Put07**] for a discussion of the issues here and for an enumeration of different reasonable definitions. Our point of view in this book is that the most important surfaces are the *closed* surfaces; however, it turns out that many of our constructions work better on surfaces with one boundary component. We view surfaces with multiple boundary components as technical tools for understanding the closed case, and we will only study them when necessary.

Disc, sphere, and torus. The Torelli groups of sufficiently simple surfaces are trivial.

LEMMA 3.2. *If Σ is either a sphere S^2, a disc $\mathbb{D}$, or a torus $\mathbb{T}^2$, then $\mathcal{I}(\Sigma) = 1$.*

PROOF. Propositions 1.1 and 1.3 say that $\text{Mod}(\mathbb{D}) = 1$ and $\text{Mod}(S^1) = 1$, so the lemma is trivial in those cases. As for the torus, Proposition 1.6 says that the action of $\text{Mod}(\mathbb{T}^2)$ on $\text{H}_1(\mathbb{T}^2; \mathbb{Z})$ is faithful, which clearly implies that $\mathcal{I}(\mathbb{T}^2) = 1$. □

One-holed torus. The first example of a surface whose Torelli group is nontrivial is a one-holed torus.

LEMMA 3.3. *We have* $\mathcal{I}_1^1 \cong \mathbb{Z}$ *with generator the Dehn twist* T_β *about the boundary component* β *of* Σ_1^1.

PROOF. Regarding Σ_1^1 as a subsurface of Σ_1 with $\Sigma_1 \backslash \Sigma_1^1$ an open disc, there is a map $\psi : \mathrm{Mod}_1^1 \to \mathrm{Mod}_1$ that extends a mapping class by the identity. Keeping in mind the equality $\mathrm{Sp}_2(\mathbb{Z}) = \mathrm{SL}_2(\mathbb{Z})$ from Remark 2.6, the symplectic representations of Mod_1^1 and Mod_1 fit into a commutative diagram

$$\begin{array}{ccc} \mathrm{Mod}_1^1 & & \\ \downarrow{\scriptstyle \psi} & \searrow & \\ \mathrm{Mod}_1 & \longrightarrow & \mathrm{SL}_2(\mathbb{Z}) \end{array}$$

Proposition 1.6 says that the map $\mathrm{Mod}_1 \to \mathrm{SL}_2(\mathbb{Z})$ is an isomorphism, so we conclude that $\mathcal{I}_1^1 = \ker(\psi)$. Theorem 1.18 implies that $\ker(\psi) \cong \mathbb{Z}$ with generator T_β. □

Higher genus. There is no simple description of $\mathcal{I}_g^b$ for $g \geqslant 2$; as we will see, it is a very large and complicated group. The only case for which we have anything like a complete description is $(g, b) = (2, 0)$. Here a theorem of Mess [**Mes92**] (see Theorem 3.14 below) says that $\mathcal{I}_2$ is an infinitely generated free group, though this theorem only gives an implicit description of the free generators. We will carefully state and prove Mess's theorem in Chapter 6.

3.2. Generators for Torelli

We now discuss several basic elements of the Torelli group.

Action of Dehn twist on homology. We begin by describing how a Dehn twist acts on homology.

LEMMA 3.4. *Let* Σ *be a surface and let* γ *be a simple closed curve on* Σ. *Orient* γ *in an arbitrary way. We then have*

$$T_\gamma(h) = h + \hat{i}([\gamma], h) \cdot [\gamma] \qquad (h \in \mathrm{H}_1(\Sigma; \mathbb{Z})). \tag{4}$$

REMARK 3.5. Equation (4) seems to depend on the orientation of γ. However, changing the orientation of γ replaces $[\gamma]$ with $-[\gamma]$ and the two minus signs in (4) cancel, so in reality (4) does not depend on the orientation.

REMARK 3.6. The map

$$h \mapsto h + \hat{i}([\gamma], h) \cdot [\gamma] \qquad (h \in \mathrm{H}_1(\Sigma; \mathbb{Z}))$$

is often called the *symplectic transvection* with respect to $[\gamma]$.

PROOF OF LEMMA 3.4. Using Theorem 2.8, it is enough to prove (4) for $h = [\delta]$, where δ is an oriented simple closed curve on Σ. Homotoping δ, we can assume that δ is immersed and transverse to γ. Equation (4) now follows from contemplating the effect of T_γ on δ. Namely, $T_\gamma(\delta)$ is obtained from δ by splicing in an arc running parallel to γ for each $p \in \gamma \cap \delta$. This arc runs in the same direction as γ if the sign of the intersection point p is positive, and this arc runs in the opposite direction of γ if the sign of the intersection point p is negative. In homology, this arc contributes $\pm[\gamma]$ with the sign the same as the sign of the intersection p. Equation (4) follows. □

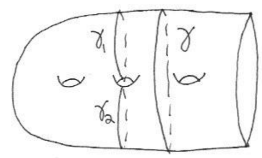

FIGURE 3.1. T_γ is a separating twist and $T_{\gamma_1} T_{\gamma_2}^{-1}$ is a bounding pair map.

Basic elements. From Lemma 3.4, we can obtain some simple elements of the Torelli group. Let Σ be a surface with at most one boundary component.

- If γ is a simple closed curve on Σ such that $[\gamma] = 0$, then Lemma 3.4 implies that T_γ acts trivially on $H_1(\Sigma; \mathbb{Z})$, i.e. that $T_\gamma \in \mathcal{I}(\Sigma)$. A simple closed curve γ satisfies $[\gamma] = 0$ if and only if γ separates Σ into two pieces (see Figure 3.1); such elements $T_\gamma \in \mathcal{I}(\Sigma)$ are therefore called *separating twists*.
- If γ_1 and γ_2 are two simple closed curves on Σ such that $[\gamma_1] = [\gamma_2]$ for some choice of orientations on the γ_i, then Lemma 3.4 implies that T_{γ_1} and T_{γ_2} act identically on $H_1(\Sigma; \mathbb{Z})$, so $T_{\gamma_1} T_{\gamma_2}^{-1} \in \mathcal{I}(\Sigma)$. The simplest such γ_i are *bounding pairs*, which are disjoint nonhomotopic nonseparating simple closed curves γ_1 and γ_2 such that $\gamma_1 \cup \gamma_2$ separates Σ into two pieces (see Figure 3.1); the associated elements $T_{\gamma_1} T_{\gamma_2}^{-1} \in \mathcal{I}(\Sigma)$ are called *bounding pair maps*.

REMARK 3.7. Though technically the boundary component β of Σ_g^1 does not separate Σ_g^1 into two pieces, we nevertheless regard T_β as a separating twist (of course, β is homotopic to a curve that separates Σ_g^1 into two pieces).

REMARK 3.8. It is easy to see that there are no bounding pairs on Σ_2.

Generating Torelli. Building on work of Birman [**Bir71**], Powell [**Pow78**] proved the following theorem.

THEOREM 3.9. *For $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, the group $\mathcal{I}_g^b$ is generated by the set of all separating twists and all bounding pair maps.*

REMARK 3.10. Observe that the generating set in Theorem 3.9 is an *infinite* generating set.

REMARK 3.11. Powell actually only explicitly deals with closed surfaces of genus at least 3, but his arguments can be easily extended to the general case.

Powell's proof of Theorem 3.9 relies on difficult combinatorial group theoretic calculations in the group $\mathrm{Sp}_{2g}(\mathbb{Z})$. More topological proofs of Theorem 3.9 were later given by Putman [**Put07**] and by Hatcher–Margalit [**HM12**]. We will give a variant of Hatcher–Margalit's proof in Chapter 7. That chapter also contains a proof of the following extension of Theorem 3.9, which is due to Johnson [**Joh79**]. The *genus* of a bounding pair $\{\gamma_1, \gamma_2\}$ on Σ_g^b with $g \geqslant 0$ and $0 \leqslant b \leqslant 1$ is defined as follows. Let S_1 and S_2 be the subsurfaces into which Σ_g^b is divided by $\gamma_1 \cup \gamma_2$. If $b = 1$, then order them so that S_2 contains the boundary component of Σ_g^b.

- If $b = 1$, then the genus of $\{\gamma_1, \gamma_2\}$ is the genus of S_1.

- If $b = 0$, then the genus of $\{\gamma_1, \gamma_2\}$ is the minimal genus of S_1 and S_2.

The genus of a bounding pair map $T_{\gamma_1} T_{\gamma_2}^{-1}$ is the genus of its underlying bounding pair.

THEOREM 3.12. *For $g \geqslant 3$ and $0 \leqslant b \leqslant 1$, the group $\mathcal{I}_g^b$ is generated by the set of all genus 1 bounding pair maps.*

REMARK 3.13. Theorem 3.12 is false for $g = 2$. Indeed, as was mentioned in Remark 3.8, there are no bounding pair maps in $\mathcal{I}_2$. Also, as we will see in §3.5, the surjection $\mathrm{Mod}_g^1 \to \mathrm{Mod}_g$ in the Birman exact sequence (see §1.4) restricts to a surjection $\mathcal{I}_g^1 \to \mathcal{I}_g$. This map takes bounding pair maps to bounding pair maps, so for $g = 2$ all bounding pair maps must be in its kernel (and thus $\mathcal{I}_2^1$ is not generated by bounding pair maps).

Finite generation. McCullough–Miller [**MM86**] proved that $\mathcal{I}_2$ is not finitely generated. As we mentioned at the end of §3.1, Mess [**Mes92**] later proved the following.

THEOREM 3.14. *The group $\mathcal{I}_2$ is an infinite rank free group.*

We will prove Theorem 3.14 in Chapter 6. We will soon show that there is a surjection $\mathcal{I}_2^1 \to \mathcal{I}_2$ (see Theorem 3.31 below), so the group $\mathcal{I}_2^1$ is not finitely generated either. We summarize this discussion in the following theorem.

THEOREM 3.15. *Neither $\mathcal{I}_2$ nor $\mathcal{I}_2^1$ is finitely generated.*

In contrast, Johnson [**Joh83**] proved the following remarkable theorem.

THEOREM 3.16. *For $g \geqslant 3$ and $0 \leqslant b \leqslant 1$, the group $\mathcal{I}_g^1$ is finitely generated.*

We will give a proof of Theorem 3.16 in Chapter 8 that combines Johnson's ideas with some later work of Putman [**Put12**].

3.3. Torsion

The main result of this section is as follows.

THEOREM 3.17. *For $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, the group $\mathcal{I}_g^b$ is torsion-free.*

REMARK 3.18. The mapping class group of a closed surface contains lots of torsion; however, the mapping class group of a surface with boundary is itself torsion-free. See Corollary 3.21 below.

The proof of Theorem 3.17 requires some preliminary results.

Realization by diffeomorphisms. The first preliminary result needed is the following theorem of Nielsen [**Nie43**].

THEOREM 3.19. *For $g, b \geqslant 0$, every finite-order element of Mod_g^b can be represented by a finite-order diffeomorphism of Σ_g^b.*

For a simple proof of Theorem 3.19 due to Fenchel [**Fen50**], see [**FM12**, §13.2]. We remark that a deep generalization of Theorem 3.19 due to Kerckhoff [**Ker83**] says that every finite subgroup of Mod_g^b can be represented by a finite group of diffeomorphisms (this is often called the Nielsen realization problem).

Isolated fixed points. The second result we need is the following lemma.

LEMMA 3.20. *Let Σ be a surface and let $F : \Sigma \to \Sigma$ be a finite-order orientation-preserving diffeomorphism of Σ such that $F \neq id$. Then all of the fixed points of f are isolated.*

PROOF. Let n be the order of F and let μ' be a Riemannian metric on Σ. Set

$$\mu = \sum_{k=0}^{n-1} F_*^{\circ k}(\mu').$$

Then μ is a Riemannian metric on Σ and F is an isometry of (Σ, μ). Let p be a fixed point of F. With respect to an orthonormal basis of the tangent space to Σ at p, the derivative of F lies in $\mathrm{SO}_2(\mathbb{R})$. If this derivative is the identity, then a standard exercise in Riemannian geometry says that $F = \mathrm{id}$; the point here is that the set of points that are fixed by F and where the derivative of F is the identity is both open and closed. Since we are assuming that $F \neq \mathrm{id}$, this cannot happen, so the derivative of F at p is not the identity. Using the fact that the derivative lies in $\mathrm{SO}_2(\mathbb{R})$, we conclude that the derivative of F at p fixes no nonzero tangent vectors, which implies that p is an isolated fixed point of F, as desired. □

COROLLARY 3.21. *For $g \geqslant 0$ and $b \geqslant 1$, the group Mod_g^b is torsion-free.*

PROOF. Let $f \in \mathrm{Mod}_g^b$ be a finite-order element. Theorem 3.19 says that we can represent f by a finite-order diffeomorphism $F : \Sigma_g^b \to \Sigma_g^b$. Since F fixes the boundary, the fixed points of F are not isolated. Lemma 3.20 then implies that $F = \mathrm{id}$, as desired. □

Finite-order diffeomorphisms and homology. The final ingredient in the proof of Theorem 3.17 is the following theorem, which was originally proved by Hurwitz [**Hur93**] in 1893 (in the same paper where he proved the Riemann-Hurwitz formula and the $84(g-1)$-theorem). The proof we will give should probably be attributed to Lefschetz [**Lef26**] (though it only appears implicitly in [**Lef26**, §71]).

THEOREM 3.22. *For $g \geqslant 2$, let $F : \Sigma_g \to \Sigma_g$ be a finite-order orientation-preserving diffeomorphism such that $F \neq id$. Then the action of F on $\mathrm{H}_1(\Sigma_g; \mathbb{Z})$ is nontrivial.*

PROOF. Lemma 3.20 says that all of the fixed points of F are isolated, so we can apply the Lefschetz fixed point theorem to it. Since F is an orientation-preserving diffeomorphism of Σ_g, the indices of all of its fixed points are 1. Letting N be the number of fixed points of F, the Lefschetz fixed point theorem implies that

$$N = \mathrm{tr}(F_* : \mathrm{H}_0(\Sigma_g; \mathbb{Q})) - \mathrm{tr}(F_* : \mathrm{H}_1(\Sigma_g; \mathbb{Q})) + \mathrm{tr}(F_* : \mathrm{H}_2(\Sigma_g; \mathbb{Q})).$$

By assumption, all of these maps are the identity. We thus conclude that $N = 2-2g$. Since $g \geqslant 2$, the quantify $2 - 2g$ is negative, a contradiction. □

REMARK 3.23. The assumption $g \geqslant 2$ in Theorem 3.22 is necessary. Indeed, for some $n \geqslant 2$ let $F_n : \mathbb{T}^2 \to \mathbb{T}^2$ be the diffeomorphism that rotates the first factor of $\mathbb{T}^2 = S^1 \times S^2$ by $2\pi/n$. Then F_n has order n, but is homotopic to the identity and thus induces the identity on $\mathrm{H}_1(\mathbb{T}^2; \mathbb{Z})$. Observe that as required by Proposition 1.6, the diffeomorphism F_n defines the trivial element in $\mathrm{Mod}(\mathbb{T}^2)$. Theorem 3.22 implies in particular that if $g \geqslant 2$, then no element of $\mathrm{Diff}_0(\Sigma_g)$ can have finite order.

Putting everything together. We close this section by proving Theorem 3.17.

PROOF OF THEOREM 3.17. Corollary 3.21 says that Mod_g^1 is torsion-free, so certainly $\mathcal{I}_g^1$ is torsion-free. It remains to show that $\mathcal{I}_g$ is torsion-free. Lemma 3.3 implies that it is enough to deal with the case $g \geqslant 2$. Consider a nonidentity torsion element $f \in \text{Mod}_g$. Theorem 3.19 says that f can be represented by a finite-order diffeomorphism $F : \Sigma_g \to \Sigma_g$. Theorem 3.22 then implies that F must act nontrivially on $\text{H}_1(\Sigma_g; \mathbb{Z})$, so $f \notin \mathcal{I}_g$, as desired. □

3.4. Action of Torelli on curves and conjugacy classes in Torelli

As we said at the beginning of §1.3, one of our main tools for studying the mapping class group is its action on the set of simple closed curves. This section is devoted to the analogue for the Torelli group of the classification of surfaces trick discussed in §1.3. This trick showed that the mapping class group acted transitively on submanifolds of the surface that "cut the surface up the same way". For the Torelli group, we will have to add appropriate homology information. The proofs here will be modeled on our proof in Chapter 2 that the symplectic representation is surjective (Theorem 2.7).

Single curves. Our first result is the following proposition (compare with Lemma 1.9). It is asserted without proof by Johnson in [**Joh80**, p. 253]; Johnson only needed a special case of it, which he proved in [**Joh80**, Lemma 10]. Our proof is modeled after [**Put07**, proof of Lemma 6.2], which proves a more precise statement that also contains Proposition 3.25 below.

PROPOSITION 3.24. *For $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, let α and α' be oriented simple closed curves on Σ_g^b. Assume that neither α nor α' separates Σ_g^b and that $[\alpha] = [\alpha']$. Then there exists some $f \in \mathcal{I}_g^b$ such that $f(\alpha) = \alpha'$.*

PROOF. Set $\alpha_1 = \alpha$ and $\alpha_1' = \alpha'$ and $a_1 = [\alpha] = [\alpha']$. Extend $\{a_1\}$ to a symplectic basis $\{a_1, b_1, \ldots, a_g, b_g\}$ for $\text{H}_1(\Sigma_g^b; \mathbb{Z})$. Proposition 2.13 implies that $\{\alpha_1\}$ and $\{\alpha_1'\}$ can be extended to geometric symplectic bases $\{\alpha_1, \beta_1, \ldots, \alpha_g, \beta_g\}$ and $\{\alpha_1', \beta_1', \ldots, \alpha_g', \beta_g'\}$ for $\text{H}_1(\Sigma_g^b; \mathbb{Z})$ such that

$$a_i = [\alpha_i] = [\alpha_i'] \quad \text{and} \quad b_i = [\beta_i] = [\beta_i']$$

for $1 \leqslant i \leqslant g$. The classification of surfaces trick then can be used to show that there exists some $f \in \text{Mod}_g^b$ such that $f(\alpha_i) = \alpha_i'$ and $f(\beta_i) = \beta_i'$ for $1 \leqslant i \leqslant g$. By construction, f fixes the basis $\{a_1, b_1, \ldots, a_g, b_g\}$ for $\text{H}_1(\Sigma_g^b; \mathbb{Z})$ pointwise, so f acts as the identity on $\text{H}_1(\Sigma_g^b; \mathbb{Z})$ and therefore lies in $\mathcal{I}_g^b$. □

Variants. The proof of Proposition 3.24 can be used to show many other similar results. For instance, it goes through without change to prove the following two propositions (compare with Lemmas 1.12 and 1.11).

PROPOSITION 3.25. *For $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, let $\{\alpha_1, \ldots, \alpha_k\}$ and $\{\alpha_1', \ldots, \alpha_k'\}$ be collections of oriented simple closed curves on Σ_g^b. Assume that the α_i are pairwise disjoint and that $\alpha_1 \cup \cdots \cup \alpha_k$ does not disconnect Σ. Similarly, assume that the α_i' are pairwise disjoint and that $\alpha_1' \cup \cdots \cup \alpha_k'$ does not disconnect Σ. Finally,* *assume that* $[\alpha_i] = [\alpha'_i]$ *for* $1 \leqslant i \leqslant k$. *Then there exists some* $f \in \mathcal{I}_g^b$ *such that* $f(\alpha_i) = \alpha'_i$ *for* $1 \leqslant i \leqslant k$.

PROPOSITION 3.26. *For* $g \geqslant 0$ *and* $0 \leqslant b \leqslant 1$, *let* $\{\alpha, \beta\}$ *and* $\{\alpha', \beta'\}$ *be collections of oriented simple closed curves on* Σ_g^b. *Assume that* α *and* β *intersect once with a positive sign. Similarly, assume that* α' *and* β' *intersect once with a positive sign. Finally, assume that* $[\alpha] = [\alpha']$ *and* $[\beta] = [\beta']$. *Then there exists some* $f \in \mathcal{I}_g^b$ *such that* $f(\alpha) = \alpha'$ *and* $f(\beta) = \beta'$.

Separating curves. For $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, let γ be a separating simple closed curve on Σ_g^b. We have $[\gamma] = 0$; however, it is definitely not the case that $\mathcal{I}_g^b$ acts transitively on separating simple closed curves (even ones that cut the surface up in the same say; see Lemma 1.13). More subtle homological information is needed.

Let S_1 and S_2 be the subsurfaces of Σ_g^b on either side of γ. If $b = 1$, then order the S_i such that S_2 contains the boundary component of Σ_g^b. Let U_1 and U_2 be the images of $\mathrm{H}_1(S_1; \mathbb{Z})$ and $\mathrm{H}_1(S_2; \mathbb{Z})$ in $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$, respectively. We then have $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z}) = U_1 \oplus U_2$. Moreover, U_1 and U_2 are symplectic subspaces of $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$ (see §2.6) which are *orthogonal* in the sense that $\hat{i}(u_1, u_2) = 0$ for all $u_1 \in U_1$ and $u_2 \in U_2$. The *separating splitting* of $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$ induced by γ is then as follows.

- If $b = 1$, then it is the ordered pair (U_1, U_2); observe that this does not depend on any choices.
- If $b = 0$, then it is the *unordered* pair (U_1, U_2). In other words, if $b = 0$ then we will identify (U_1, U_2) and (U_2, U_1). This is required since the S_i are not canonically ordered.

With these definitions, we have the following.

PROPOSITION 3.27. *For* $g \geqslant 0$ *and* $0 \leqslant b \leqslant 1$, *let* γ *and* γ' *be simple closed separating curves on* Σ_g^b. *Then there exists some* $f \in \mathcal{I}_g^b$ *such that* $f(\gamma) = \gamma'$ *if and only if* γ *and* γ' *induce the same separating splitting of* $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$.

PROOF. It is clear that γ and γ' induce the same separating splitting if such an f exists. Conversely, assume that γ and γ' induce the same separating splitting of $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$. Let S_1 and S_2 (resp. S'_1 and S'_2) be the subsurfaces of Σ_g^b on either side of γ (resp. γ'). If $b = 1$, then order the S_i and S'_i such that S_2 and S'_2 contain the boundary component of Σ_g^b; by assumption, the images of $\mathrm{H}_1(S_i; \mathbb{Z})$ and $\mathrm{H}_1(S'_i; \mathbb{Z})$ in $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$ are then the same for $i = 1, 2$. If $b = 0$, then simply order the S_i and S'_i such that the images of $\mathrm{H}_1(S_i; \mathbb{Z})$ and $\mathrm{H}_1(S'_i; \mathbb{Z})$ in $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$ are the same for $i = 1, 2$. Observe that these choices of orderings implies that $S_i \cong S'_i$ for $i = 1, 2$. For $i = 1, 2$, let U_i be the common image of $\mathrm{H}_1(S_i; \mathbb{Z})$ and $\mathrm{H}_1(S'_i; \mathbb{Z})$ in $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$, so $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z}) = U_1 \oplus U_2$.

Let h be the genus of S_1; by construction, h is also the genus of S'_1. Both S_2 and S'_2 then have genus $g - h$. Let $\{a_1, b_1, \ldots, a_h, b_h\}$ be a symplectic basis for U_1 and let $\{a_{h+1}, b_{h+1}, \ldots, a_g, b_g\}$ be a symplectic basis for U_2. Applying Proposition 2.12 twice, we can find a geometric symplectic basis $\{\alpha_1, \beta_1, \ldots, \alpha_g, \beta_g\}$ on Σ_g^b such that

$$\alpha_1, \beta_1, \ldots, \alpha_h, \beta_h \subset S_1 \quad \text{and} \quad \alpha_{h+1}, \beta_{h+1}, \ldots, \alpha_g, \beta_g \subset S_2$$

and such that

$$[\alpha_i] = a_i \quad \text{and} \quad [\beta_i] = b_i \qquad (1 \leqslant i \leqslant g).$$

Similarly, we can find a geometric symplectic basis $\{\alpha'_1, \beta'_1, \ldots, \alpha'_g, \beta'_g\}$ on Σ_g^b such that

$$\alpha'_1, \beta'_1, \ldots, \alpha'_h, \beta'_h \subset S'_1 \quad \text{and} \quad \alpha'_{h+1}, \beta'_{h+1}, \ldots, \alpha'_g, \beta'_g \subset S'_2$$

and such that

$$[\alpha_i] = a_i \quad \text{and} \quad [\beta_i] = b_i \qquad (1 \leqslant i \leqslant g).$$

Using the classification of surfaces trick (see §1.3), we can find $f \in \mathrm{Mod}_g^b$ such that $f(\gamma) = \gamma'$ and such that

$$f(\alpha_i) = \alpha'_i \quad \text{and} \quad f(\beta_i) = \beta'_i \qquad (1 \leqslant i \leqslant g).$$

By construction, f fixes a_i and b_i for $1 \leqslant i \leqslant g$, so $f \in \mathcal{I}_g^b$. □

This has the following corollary, which was originally proven by Johnson [**Joh80**].

COROLLARY 3.28. *For $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, let T_γ and $T_{\gamma'}$ be separating twists on Σ_g^b. Then T_γ and $T_{\gamma'}$ are conjugate in $\mathcal{I}_g^b$ if and only if γ and γ' induce the same separating splitting of $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$.*

PROOF. The separating twists T_γ and $T_{\gamma'}$ are conjugate in $\mathcal{I}_g^b$ if and only if there exists some $f \in \mathcal{I}_g^b$ such that $fT_\gamma f^{-1} = T_{\gamma'}$. Since $fT_\gamma f^{-1} = T_{f(\gamma)}$ (see Lemma 1.23), this holds if and only if there exists some $f \in \mathcal{I}_g^b$ such that $f(\gamma) = \gamma'$. Proposition 3.27 says that this holds if and only if γ and γ' induce the same separating splitting of $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$. Here we are using the fact that two Dehn twists are equal if and only if their associated curves are homotopic; see [**FM12**, Fact 3.6]. □

Bounding pairs. We now turn to the action of the Torelli group on bounding pairs. Again, the necessary invariant will be the homology of the subsurfaces on either side of the bounding pair. This will require the following two pieces of notation. First, if M is a free finite-rank $\mathbb{Z}$-module equipped with a symplectic form $\omega(\cdot, \cdot)$ and $x \in M$, then define $x^\perp = \{y \in M \mid \omega(x, y) = 0\}$. Second, if η is an oriented simple closed curve on a surface, then $\overline{\eta}$ is η with the opposite orientation.

For $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, let $\{\gamma_1, \gamma_2\}$ be a bounding pair. If $b = 1$, then orient γ_1 and γ_2 such that the boundary component of Σ_g^b is in the subsurface of Σ_g^b to the right of $\gamma_1 \cup \overline{\gamma}_2$; this will ensure that $[\gamma_1] = [\gamma_2]$. If $b = 0$, then orient γ_1 and γ_2 arbitrarily such that $[\gamma_1] = [\gamma_2]$. Let S_1 be the subsurface of Σ_g^b to the left of $\gamma_1 \cup \overline{\gamma}_2$ and S_2 be the subsurface of Σ_g^b to their right. For $i = 1, 2$ let U_i be the image of $\mathrm{H}_1(S_i; \mathbb{Z})$ in $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$. Clearly U_1 and U_2 together span $[\gamma_1]^\perp$ and $U_1 \cap U_2 = \langle [\gamma_1] \rangle$. The *BP-splitting* induced by $\{\gamma_1, \gamma_2\}$ is as follows.

- If $b = 1$, then it is the ordered triple $([\gamma_1], U_1, U_2)$. Observe that this does not depend on any choices.
- If $b = 0$, then it is the ordered triple $([\gamma_1], U_1, U_2)$. However, since this depends on a choice of orientation for γ_1, we will identify $([\gamma_1], U_1, U_2)$ and $(-[\gamma_1], U_2, U_1)$. With these identifications, this is well-defined.

With these definitions, we have the following.

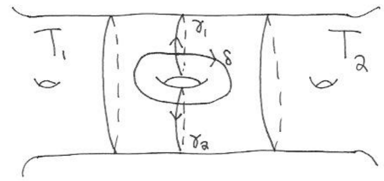

FIGURE 3.2. The bounding pair $\{\gamma_1, \gamma_2\}$ and the curve δ that intersects both γ_1 and γ_2 once. The surface S_1 is the subsurface of Σ_g^b on the left of $\gamma_1 \cup \overline{\gamma}_2$ and T_1 is the complement in S_1 of a regular neighborhood of $\gamma_1 \cup (\delta \cap S_1) \cup \gamma_2$. Similarly, S_2 is the subsurface of Σ_g^b on the right of $\gamma_1 \cup \overline{\gamma}_2$ and T_2 is the complement in S_2 of a regular neighborhood of $\gamma_1 \cup (\delta \cap S_2) \cup \gamma_2$.

PROPOSITION 3.29. *For $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, let $\{\gamma_1, \gamma_2\}$ and $\{\gamma_1', \gamma_2'\}$ be bounding pairs on Σ_g^b. Then there exists some $f \in \mathcal{I}_g^b$ such that $f(\gamma_i) = \gamma_i'$ for $i = 1, 2$ if and only if $\{\gamma_1, \gamma_2\}$ and $\{\gamma_1', \gamma_2'\}$ induce the same BP-splitting.*

PROOF. It is clear that $\{\gamma_1, \gamma_2\}$ and $\{\gamma_1', \gamma_2'\}$ induce the same BP-splitting if such an f exists. Conversely, assume that $\{\gamma_1, \gamma_2\}$ and $\{\gamma_1', \gamma_2'\}$ induce the same BP-splitting. Orient the γ_i as in the definition of a BP-splitting, so $[\gamma_1] = [\gamma_2]$ and if $b = 1$ then the boundary component of Σ_g^b lies in the subsurface of Σ_g^b to the right of $\gamma_1 \cup \overline{\gamma}_2$. Orient the γ_i' in a similar way; if $b = 0$, then make the choice of orientation that assures that $[\gamma_1] = [\gamma_1']$ (if $b = 1$, then this is automatic).

Let the common BP-splitting of $\{\gamma_1, \gamma_2\}$ and $\{\gamma_1', \gamma_2'\}$ be $([\gamma_1], U_1, U_2)$. Let S_1 and S_2 (resp. S_1' and S_2') be the subsurfaces of Σ_g^b to the left and to the right of $\gamma_1 \cup \overline{\gamma}_2$ (resp. $\gamma_1' \cup \overline{\gamma}_2'$), respectively. By definition, for $i = 1, 2$ the subspace U_i of $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$ is the image of both $\mathrm{H}_1(S_i; \mathbb{Z})$ and $\mathrm{H}_1(S_i'; \mathbb{Z})$ in $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$.

Let $d \in \mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$ be such that $\hat{i}(d, [\gamma_1]) = \hat{i}(d, [\gamma_1']) = 1$. Using Lemma 2.15, we can find oriented simple closed curves δ and δ' such that $[\delta] = [\delta'] = d$ and such that δ (resp. δ') intersects both γ_1 and γ_2 (resp. γ_1' and γ_2') once. For $i = 1, 2$ define T_i to be the complement in S_i of an open regular neighborhood of $\gamma_1 \cup (\delta \cap S_i) \cup \gamma_2$; see Figure 3.2. Similarly, for $i = 1, 2$ define T_i' to be the complement in S_i' of an open regular neighborhood of $\gamma_1' \cup (\delta' \cap S_i') \cup \gamma_2$. Observe that all of the boundary components of T_i and T_i' are separating curves in Σ_g^b. As we discussed in §2.6, this implies that the image V_i (resp. V_i') of $\mathrm{H}_1(T_i; \mathbb{Z})$ (resp. $\mathrm{H}_1(T_i'; \mathbb{Z})$) in $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$ is a symplectic subspace.

The key to our proof now is the observation that for $i = 1, 2$ we have $V_i = V_i'$; indeed, these are both equal to $\{x \in U_i \mid \hat{i}(x, d) = 0\}$. Moreover, as is clear from Figure 3.2, if T_1 and T_1' have genus h, then T_2 and T_2' have genus $g - h - 1$. Let $\{a_1, b_1, \ldots, a_h, b_h\}$ be a symplectic basis for $V_1 = V_1'$ and let $\{a_{h+2}, b_{h+2}, \ldots, a_g, b_g\}$ be a symplectic basis for $V_2 = V_2'$; the set

$$\{a_1, b_1, \ldots, a_h, b_h, d, [\gamma_1], a_{h+2}, b_{h+2}, \ldots, a_g, b_g\}$$

is then a symplectic basis for $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$. Using Proposition 2.12, we can find the following.

- Simple closed curves $\{\alpha_1, \beta_1, \ldots, \alpha_h, \beta_h\}$ in T_1 such that $[\alpha_i] = a_i$ and $[\beta_i] = b_i$ for $1 \leqslant i \leqslant h$.

- Simple closed curves $\{\alpha'_1, \beta'_1, \ldots, \alpha'_h, \beta'_h\}$ in T'_1 such that $[\alpha'_i] = a_i$ and $[\beta'_i] = b_i$ for $1 \leqslant i \leqslant h$.
- Simple closed curves $\{\alpha_{h+2}, \beta_{h+2}, \ldots, \alpha_g, \beta_g\}$ in T_2 such that $[\alpha_i] = a_i$ and $[\beta_i] = b_i$ for $h+2 \leqslant i \leqslant g$.
- Simple closed curves $\{\alpha'_{h+2}, \beta'_{h+2}, \ldots, \alpha'_g, \beta'_g\}$ in T'_2 such that $[\alpha'_i] = a_i$ and $[\beta'_i] = b_i$ for $h+2 \leqslant i \leqslant g$.

The intersection patterns of these curves will be such that the sets

$$\{\alpha_1, \beta_1, \ldots, \alpha_h, \beta_h, \delta, \gamma_1, \alpha_{h+2}, \beta_{h+2}, \ldots, \alpha_g, \beta_g\}$$

and

$$\{\alpha'_1, \beta'_1, \ldots, \alpha'_h, \beta'_h, \delta', \gamma'_1, \alpha'_{h+2}, \beta'_{h+2}, \ldots, \alpha'_g, \beta'_g\}$$

are both geometric symplectic bases for $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$. The classification of surfaces trick then implies that there exists some $f \in \mathrm{Mod}_g^b$ such that $f(\alpha_i) = \alpha'_i$ and $f(\beta_i) = \beta'_i$ for $1 \leqslant i \leqslant g$ with $i \neq h+1$ and such that

$$f(\delta) = \delta' \quad \text{and} \quad f(\gamma_1) = \gamma'_1 \quad \text{and} \quad f(\gamma_2) = \gamma'_2.$$

Since by construction f pointwise fixes a symplectic basis for $\mathrm{H}_1(\Sigma_g^b; \mathbb{Z})$, we have that $f \in \mathcal{I}_g^b$, and we are done. □

This has the following corollary, which was originally proven by Johnson [**Joh80**].

Corollary 3.30. *For $g \geqslant 0$ and $0 \leqslant b \leqslant 1$, let $T_{\gamma_1} T_{\gamma_2}^{-1}$ and $T_{\gamma_1} T_{\gamma_2}^{-1}$ be bounding pair maps on Σ_g^b. Then $T_{\gamma_1} T_{\gamma_2}^{-1}$ and $T_{\gamma'_1} T_{\gamma'_2}^{-1}$ are conjugate in $\mathcal{I}_g^b$ if and only if $\{\gamma_1, \gamma_2\}$ and $\{\gamma'_1, \gamma'_2\}$ induce the same BP-splitting.*

Proof. The bounding pair maps $T_{\gamma_1} T_{\gamma_2}^{-1}$ and $T_{\gamma'_1} T_{\gamma'_2}^{-1}$ are conjugate in $\mathcal{I}_g^b$ if and only if there exists some $f \in \mathcal{I}_g^b$ such that

$$f T_{\gamma_1} T_{\gamma_2}^{-1} f^{-1} = T_{\gamma'_1} T_{\gamma'_2}^{-1}.$$

Since

$$f T_{\gamma_1} T_{\gamma_2}^{-1} f^{-1} = T_{f(\gamma_1)} T_{f(\gamma_2)}^{-1},$$

this holds if and only if there exists some $f \in \mathcal{I}_g^b$ such that $f(\gamma_1) = \gamma'_1$ and $f(\gamma_2) = \gamma'_2$. Proposition 3.27 says that this holds if and only if $\{\gamma_1, \gamma_2\}$ and $\{\gamma'_1, \gamma'_2\}$ induce the same BP-splitting. □

3.5. Closed surfaces vs surfaces with boundary

We close this chapter by discussing the relationship between the Torelli groups on closed surfaces and on surfaces with boundary. Recall that for the mapping class group, this is given by the Birman exact sequence (see Theorem 1.17). There is a similar result for the Torelli group.

Theorem 3.31. *For $g \geqslant 2$, there is a short exact sequence*

$$1 \longrightarrow \pi_1(U\Sigma_g) \longrightarrow \mathcal{I}_g^1 \longrightarrow \mathcal{I}_g \longrightarrow 1.$$

Proof. Our proof builds on the proof of Lemma 3.3 above. Theorem 1.17 says that there is a short exact sequence

$$1 \longrightarrow \pi_1(U\Sigma_g) \longrightarrow \mathrm{Mod}_g^1 \xrightarrow{\psi} \mathrm{Mod}_g \longrightarrow 1. \tag{5}$$

Regarding Σ_g^1 as a subsurface of Σ_g with $\Sigma_g \backslash \Sigma_g^1$ an open disc, recall that the map $\psi : \mathrm{Mod}_g^1 \to \mathrm{Mod}_g$ in (5) is the map that extends mapping classes by the identity. The symplectic representations of Mod_g^1 and Mod_g fit into a commutative diagram

$$\begin{array}{ccc} \mathrm{Mod}_g^1 & & \\ \downarrow{\scriptstyle \psi} & \searrow & \\ \mathrm{Mod}_g & \longrightarrow & \mathrm{Sp}_{2g}(\mathbb{Z}) \end{array}$$

It follows this that $\ker(\psi) = \pi_1(U\Sigma_g)$ is contained in $\mathcal{I}_g^1$. Since ψ is surjective, it also follows that the image of $\psi|_{\mathcal{I}_g^1}$ is $\mathcal{I}_g$. Summing up, (5) restricts to a short exact sequence

$$1 \longrightarrow \pi_1(U\Sigma_g) \longrightarrow \mathcal{I}_g^1 \longrightarrow \mathcal{I}_g \longrightarrow 1,$$

as desired. □

Part 2

Combinatorial Group Theory

Introduction to Part 2

In Chapter 3, we stated without proof the following four fundamental theorems about the combinatorial group theory of the Torelli group.

- Birman and Powell's theorem [**Bir71, Pow78**] asserting that the Torelli group is generated by separating twists and bounding pair maps (Theorem 3.9).
- Johnson's theorem [**Joh79**] asserting that in genus at least 3 the Torelli group is generated by bounding pair maps (Theorem 3.12).
- Mess's theorem [**Mes92**] asserting that the genus 2 Torelli group is an infinite rank free group (Theorem 3.14).
- Johnson's theorem [**Joh83**] asserting that the Torelli group is finitely generated when the genus is at least 3 (Theorem 3.16).

This part of the book is devoted to the proofs of these four results.

Results needed later. A reader encountering the Torelli group for the first time might want to take the results of this part of the book on faith and skip immediately to Parts 3 and 4, which cover the Johnson homomorphism and connections to 3-manifolds. In fact, for the vast majority of Parts 3 and 4 the only result needed from Part 2 is the fact that the Torelli group is generated by separating twists and bounding pair maps. There are only two exceptions to this.

- Chapter 13 of Part 3 is devoted to proving a theorem of Johnson [**Joh85a**] which says that the subgroup of the Torelli group generated by separating twists is exactly the kernel of the Johnson homomorphism. This needs two additional results from Part 2. The first is the *decomposition theorem*, which is a technical result proved in Chapter 7 that describes how the genus g Torelli group is "built" from lower-genus pieces. The second is a theorem from Chapter 5 which describes the stabilizer in the Torelli group of a nonseparating simple closed curve.
- Chapter 18 of Part 4 is devoted to proving a theorem of Johnson [**Joh85b**] which gives the abelianization of the Torelli group. This chapter has the same set of prerequisites as Chapter 13.

Outline. The outline of this part of the book is as follows. We begin with two technical chapters. Chapter 4 discusses the *complex of reduced cycles*, which is a space encoding all the ways that a fixed homology class can be written as a cycle on the surface, and Chapter 5 describes the stabilizer in the Torelli group of a nonseparating simple closed curve. We then prove Theorem 3.14 in Chapters 6, Theorems 3.9 and 3.12 in Chapter 7 (which also contains the decomposition theorem), and Theorem 3.16 in Chapter 8.

CHAPTER 4

The complex of cycles

This chapter is devoted to the *complex of reduced cycles* on a surface, which is a space that encodes all the ways that a fixed homology class can be written as a cycle. This complex was introduced by Bestvina–Bux–Margalit [**BBM10**], who used it to calculate the cohomological dimension of the Torelli group and to give a topological proof of a theorem of Mess [**Mes92**] that says that $\mathcal{I}_2$ is an infinitely generated free group. We will give this proof of Mess's theorem in Chapter 6. We will also use the complex of reduced cycles to prove the *decomposition theorem* in Chapter 7, which is basic structural result about the Torelli group that we will use to prove many things.

The main result in this chapter is Theorem 4.14, which asserts that the complex of reduced cycles is contractible. Our proof follows the "second proof" of this fact from [**BBM10**]. See [**HM12**] for an alternate exposition of it.

Throughout this section, Σ is a closed surface and $x \in \mathrm{H}_1(\Sigma; \mathbb{Z})$ is a fixed primitive element.

4.1. Basic definitions

We will first define the complex of cycles as a set and then discuss its topology.

Multicurves and weighted multicurves. An *oriented multicurve* γ on Σ is an unordered collection $\gamma_1 \cup \cdots \cup \gamma_k$ of disjoint oriented nonnullhomotopic simple closed curves which are pairwise non-homotopic (as unoriented curves, i.e. we do not allow one of the γ_i to be homotopic to another γ_j but with a reversed orientation). We will not distinguish between homotopic oriented multicurves. A *submulticurve* of an oriented multicurve γ is an oriented multicurve each of whose curves is also a curve in γ.

A *weighted oriented multicurve* is a formal expression $c_1\gamma_1 + \cdots + c_k\gamma_k$ with $\gamma := \gamma_1 \cup \cdots \cup \gamma_k$ an oriented multicurve and $c_1, \ldots, c_k \in \mathbb{R}$. The ordering of the γ_i in this expression does not matter. The number c_i is the *weight* of γ_i and the *support* of $c_1\gamma_1 + \cdots + c_k\gamma_k$ is the submulticurve of γ composed of all of the γ_i whose weights are nonzero. We will identify two weighted oriented multicurves that differ by inserting or deleting oriented curves of weight 0. The homology class *represented* by $c_1\gamma_1 + \cdots + c_k\gamma_k$ is $c_1[\gamma_1] + \cdots + c_k[\gamma_k] \in \mathrm{H}_1(\Sigma; \mathbb{Z})$. We will call $c_1\gamma_1 + \cdots + c_k\gamma_k$ a *positively weighted oriented multicurve* if all the weights c_i are nonnegative.

Complex of unreduced cycles as a set. The *complex of unreduced cycles*, denoted $\widehat{\mathcal{C}}_x(\Sigma)$, is the set of positively weighted oriented multicurves representing the fixed primitive homology class x. We will soon define a topology on $\widehat{\mathcal{C}}_x(\Sigma)$. Intuitively, one moves around in this topology by continuously varying the weights

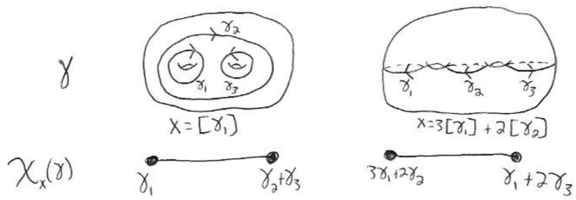

FIGURE 4.1. Example of multicurves γ whose associated cells $\mathcal{X}_x(\gamma)$ are compact 1-dimensional polyhedra. Under each multicurve is how to write x as a linear combination of the multicurves. The points in the interior of the left edge are $t\gamma_1 + s\gamma_2 + s\gamma_2$ with $s, t \geqslant 0$ and $t + s = 1$. The points in the interior of the right edge are $(3t + s)\gamma_1 + 2t\gamma_2 + (2t + 2s)\gamma_3$ with $s, t \geqslant 0$ and $s + t = 1$.

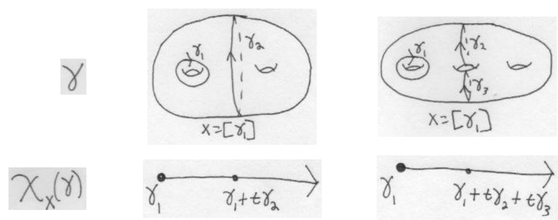

FIGURE 4.2. Examples of multicurves γ whose associated cells $\mathcal{X}_x(\gamma)$ are noncompact 1-dimensional polyhedra. Under each multicurve is how to write x as a linear combination of the multicurves.

in positively weighted oriented multicurves while keeping the represented homology class constant. When one of the weights goes to 0, that curve disappears.

Cells. Let γ be some oriented multicurve on Σ. The *cell* associated to γ, denoted $\mathcal{X}_x(\gamma)$, is the subset of $\widehat{\mathcal{C}}_x(\Sigma)$ consisting of positively weighted oriented multicurves representing x whose support is a submulticurve of γ. We will say that $\mathcal{X}_x(\gamma)$ is *nondegenerate* if it contains a positively weighted oriented multicurve whose support is equal to γ. Writing $\gamma = \gamma_1 \cup \cdots \cup \gamma_k$, there is an inclusion of $\mathcal{X}_x(\gamma)$ into $\mathbb{R}^k_{\geqslant 0}$ that takes $c_1\gamma_1 + \cdots + c_k\gamma_k$ to $(c_1, \ldots, c_k)$. This inclusion defines a topology on $\mathcal{X}_x(\gamma)$; in fact, it endows $\mathcal{X}_x(\gamma)$ with the structure of a (not necessarily compact) polyhedron, possibly empty. This structure does not depend the ordering of the γ_i. If γ' is a submulticurve of γ, then $\mathcal{X}_x(\gamma')$ is in a natural way a subpolyhedron of $\mathcal{X}_x(\gamma)$. See Figures 4.1–4.3 for some examples of cells.

Topology on complex of unreduced cycles. If γ is an oriented multicurve on Σ, then $\mathcal{X}_x(\gamma)$ can be regarded as a subset of $\widehat{\mathcal{C}}_x(\Sigma)$. We will give $\widehat{\mathcal{C}}_x(\Sigma)$ the weak topology with regards to the $\mathcal{X}_x(\gamma)$. In other words, a set $U \subset \widehat{\mathcal{C}}_x(\Sigma)$ is open if and only if $U \cap \mathcal{X}_x(\gamma)$ is open for all oriented multicurves γ.

Complex of reduced cycles. We will say that a cell $\mathcal{X}_x(\gamma)$ is *reduced* if it is compact. Below in Lemma 4.4 we will give an easy-to-check characterization of when a cell is reduced. We will call a positively weighted oriented multicurve $c \in \widehat{\mathcal{C}}_x(\Sigma)$ with support γ *reduced* if $\mathcal{X}_x(\gamma)$ is reduced. The *complex of reduced cycles*,

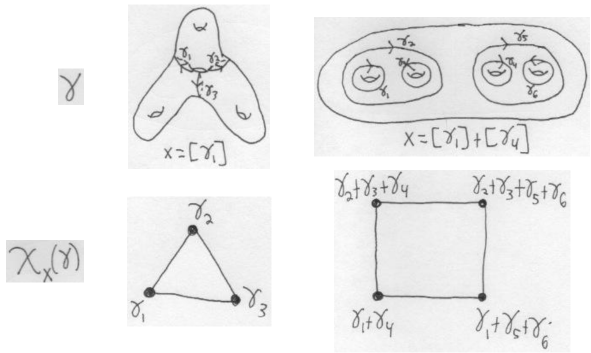

FIGURE 4.3. Examples of multicurves γ whose associated cells $\mathcal{X}_x(\gamma)$ are compact 2-dimensional polyhedra. Under each multicurve is how to write x as a linear combination of the multicurves. The points in the interior of the triangle are $t_1\gamma_1 + t_2\gamma_2 + t_3\gamma_3$ with $t_1, t_2, t_3 \geqslant 0$ and $t_1 + t_2 + t_3 = 1$. The points in the interior of the square are $t\gamma_1 + s\gamma_2 + s\gamma_3 + t'\gamma_4 + s'\gamma_5 + s'\gamma_6$ with $t, t', s, s' \geqslant 0$ and $t + s = 1$ and $t' + s' = 1$. We remark that one can also find cells that are pentagons, hexagons, etc.

denoted $\mathcal{C}_x(\Sigma)$, is the subset of $\widehat{\mathcal{C}}_x(\Sigma)$ consisting of reduced positively weighted oriented multicurves representing x. The reduced cells endow $\mathcal{C}_x(\Sigma)$ with the structure of a polyhedral complex.

4.2. Basic properties of cells

We now discuss some basic properties of cells. Throughout this section, $\gamma = \gamma_1 \cup \cdots \cup \gamma_k$ is a fixed oriented multicurve on Σ such that $\mathcal{X}_x(\gamma)$ is nondegenerate.

Zero sets. Define $\mathcal{Z}(\gamma)$ to be the set of all weighted oriented multicurves $c_1\gamma_1 + \cdots + c_k\gamma_k$ that represent $0 \in \mathrm{H}_1(\Sigma; \mathbb{Z})$. Just like for $\mathcal{X}_x(\gamma)$, we can identify $\mathcal{Z}(\gamma)$ with a subset (in fact, a linear subspace) of $\mathbb{R}^k$. Since we are assuming that $\mathcal{X}_x(\gamma)$ is nondegenerate, under these identifications $\mathcal{X}_x(\gamma)$ is the intersection of an affine subset of $\mathbb{R}^k$ parallel to $\mathcal{Z}(\gamma)$ with the positive orthant $\mathbb{R}^k_{\geqslant 0}$.

Generators and relations for zero sets. Let R be a subsurface of Σ whose boundary components (considered as *unoriented* curves) lie in γ. Letting $1 \leqslant i_1 < i_2 < \cdots < i_p \leqslant k$ be the indices such that the boundary components of R are $\gamma_{i_1}, \ldots, \gamma_{i_p}$, define

$$\partial R = \pm\gamma_{i_1} + \cdots + \pm\gamma_{i_p},$$

where the signs reflect whether or not the orientation of γ_{i_j} agrees or not with the orientation it acquires from R. Clearly $\partial R \in \mathcal{Z}(\gamma)$. We then have the following.

LEMMA 4.1. *Let γ be an oriented multicurve on Σ. Let $R_1, \ldots, R_\ell$ be the connected subsurfaces of Σ obtained by cutting Σ along γ. Then $\mathcal{Z}(\gamma)$ is generated by $\{\partial R_1, \ldots, \partial R_\ell\}$, and the only relation between these generators is $\partial R_1 + \cdots + \partial R_\ell =$* 0.

This lemma has the following corollary.

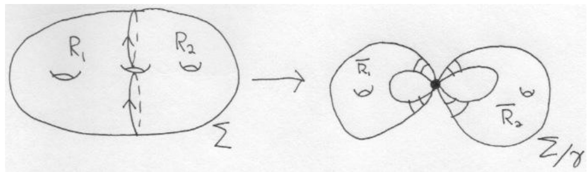

FIGURE 4.4. The result Σ/γ of collapsing a multicurve to a point. For $i = 1, 2$, the surface $\overline{R}_i$ can both be obtained from a torus by identifying two of its points together.

COROLLARY 4.2. *Let γ be an oriented multicurve on Σ such that $\mathcal{X}_x(\gamma)$ is nondegenerate. Let $\ell \geqslant 1$ be the number of components of Σ cut along γ. Then $\mathcal{X}_x(\gamma)$ is an $(\ell - 1)$-dimensional (not necessarily compact) polyhedron.*

PROOF OF LEMMA 4.1. We can identify $\mathcal{Z}(\gamma)$ with the kernel of the map $\mathrm{H}_1(\gamma; \mathbb{Z}) \to \mathrm{H}_1(\Sigma; \mathbb{Z})$. The long exact sequence in homology associated to the pair (Σ, γ) therefore induces an exact sequence

$$\mathrm{H}_2(\Sigma; \mathbb{Z}) \longrightarrow \mathrm{H}_2(\Sigma/\gamma; \mathbb{Z}) \xrightarrow{\pi} \mathcal{Z}(\gamma) \longrightarrow 0.$$

Letting $\overline{R}_i$ be the image of R_i in Σ/γ (see Figure 4.4), we have an element $[\overline{R}_i] \in \mathrm{H}_2(\Sigma/\gamma; \mathbb{Z})$ satisfying $\pi([\overline{R}_i]) = \partial R_i$. The group $\mathrm{H}_2(\Sigma/\gamma; \mathbb{Z})$ is the free abelian group with basis $\{[\overline{R}_1], \ldots, [\overline{R}_\ell]\}$, and the generator $[\Sigma]$ of $\mathrm{H}_2(\Sigma; \mathbb{Z}) \cong \mathbb{Z}$ maps to $[\overline{R}_1] + \cdots + [\overline{R}_\ell]$. The lemma follows. □

Vertices. We now give a concrete description of the vertices of $\widehat{\mathcal{C}}_x(\Sigma)$ (which of course coincide with the vertices of $\mathcal{C}_x(\Sigma)$).

LEMMA 4.3. *The vertices of $\widehat{\mathcal{C}}_x(\Sigma)$ consist of $c_1\gamma_1 + \cdots + c_k\gamma_k$, where $\gamma_1 \cup \cdots \cup \gamma_k$ is an oriented multicurve on Σ that does not separate Σ and the c_i are positive integers such that*

$$c_1[\gamma_1] + \cdots + \cdots c_k[\gamma_k] = x.$$

PROOF. Consider a point $c = c_1\gamma_1 + \cdots + c_k\gamma_k$ of $\widehat{\mathcal{C}}_x(\Sigma)$. Assume that none of the c_i vanish, so the support of c is $\gamma = \gamma_1 \cup \cdots \cup \gamma_k$. The cell $\mathcal{X}_x(\gamma)$ is therefore nondegenerate. The point c is a vertex of $\widehat{\mathcal{C}}_x(\Sigma)$ exactly when $\mathcal{X}_x(\gamma)$ is 0-dimensional (and hence consists of the single point c). Corollary 4.2 says that this holds if and only if $\gamma_1 \cup \cdots \cup \gamma_k$ does not separate Σ. It remains to prove that if c is a vertex, then each c_i is an integer.

Consider some $1 \leqslant i \leqslant k$. Since $\gamma_1 \cup \cdots \cup \gamma_k$ does not separate Σ, we can find an oriented simple closed curve δ on Σ that intersects γ_i once with a positive sign and is disjoint from γ_j for $1 \leqslant j \leqslant k$ with $j \neq i$. We then have

$$c_i = \hat{i}([\delta], c_1[\gamma_1] + \cdots + c_k[\gamma_k]) = \hat{i}([\delta], x) \in \mathbb{Z};$$

the final inequality follows from the fact that $x \in \mathrm{H}_1(\Sigma; \mathbb{Z})$. □

Criterion for being reduced. We now prove the following simple description of when a cell is reduced (compare with the examples in Figure 4.2).

LEMMA 4.4. *Let $\gamma = \gamma_1 \cup \cdots \cup \gamma_k$ be a multicurve such that that $\mathcal{X}_x(\gamma)$ is nondegenerate. The cell $\mathcal{X}_x(\gamma)$ is reduced if and only if there does not exist $1 \leqslant i_1 < \cdots < i_p \leqslant k$ such that $[\gamma_{i_1}] + \cdots + [\gamma_{i_p}] = 0$.*

PROOF. If there exist $1 \leqslant i_1 < \cdots < i_p \leqslant k$ such that $[\gamma_{i_1}] + \cdots + [\gamma_{i_p}] = 0$, then fixing some point $c \in \mathcal{X}_x(\gamma)$ we have an infinite ray

$$\{c + t(\gamma_{i_1} + \cdots + \gamma_{i_p}) \mid t \geqslant 0\} \subset \mathcal{X}_x(\gamma).$$

Thus $\mathcal{X}_x(\gamma)$ is noncompact, and hence not reduced.

We will prove the contrapositive of the other implication of the lemma. Assume that $\mathcal{X}_x(\gamma)$ is nonreduced, i.e. not compact. We first prove that there exist real numbers $c_1, \ldots, c_k \geqslant 0$ (not all 0) such that $c_1[\gamma_1] + \cdots + c_k[\gamma_k] = 0$. Since $\mathcal{X}_x(\gamma)$ is a noncompact polyhedron, it must contain an infinite ray. Let $c' = c'_1\gamma_1 + \cdots + c'_k\gamma_k$ be the initial point of this ray and let $c'' = c''_1\gamma_1 + \cdots + c''_k\gamma_k$ be some other point on this ray. For $1 \leqslant i \leqslant k$, set $c_i = c''_i - c'_i$. We thus have

$$c_1[\gamma_1] + \cdots + c_k[\gamma_k] = (c'_1[\gamma_1] + \cdots + c'_k[\gamma_k]) - (c''_1[\gamma_1] + \cdots + c''_k[\gamma_k]) = x - x = 0.$$

Moreover, the points

$$\{(c'_1 + tc_1)\gamma_1 + \cdots + (c'_k + tc_k)\gamma_k \mid t \geqslant 0\}$$

all lie in $\mathcal{X}_x(\gamma)$, i.e. $c'_i + tc_i \geqslant 0$ for all $t \geqslant 0$ and all $1 \leqslant i \leqslant k$. We conclude that $c_i \geqslant 0$, as desired.

Let $R_1, \ldots, R_\ell$ be the connected subsurfaces of Σ obtained by cutting Σ along γ. Lemma 4.1 implies that there exists some $d_1, \ldots, d_\ell \in \mathbb{R}$ such that

$$c_1\gamma_1 + \cdots + c_k\gamma_k = d_1\partial R_1 + \cdots d_\ell\partial R_\ell.$$

Since $\partial R_1 + \cdots + \partial R_\ell = 0$, we can add a large positive constant E to each d_i and ensure that $d_i > 0$ for $1 \leqslant i \leqslant \ell$. Set $d = \max\{d_1, \ldots, d_\ell\}$, and assume that the R_j are ordered such that $d_1 = \cdots = d_r = d$ and $d_{r+1}, \ldots, d_\ell < d$. Since not all the c_i are 0 and $\partial R_1 + \cdots + \partial R_\ell = 0$, we must have $r < \ell$. Setting $R = R_1 \cup \cdots \cup R_r$, the surface R is thus a proper subsurface of Σ, so $\partial R \neq 0$. Observe that

$$c_1\gamma_1 + \cdots + c_k\gamma_k = d\partial R + d_{r+1}\partial R_{r+1} + \cdots + d_\ell\partial R_\ell.$$

Each γ_i occurs as the boundary of exactly two of the R_j. Since $d_j < d$ for $r+1 \leqslant j \leqslant \ell$ and $c_i \geqslant 0$ for all $1 \leqslant i \leqslant k$, the coefficients of all of the γ_i which appear in ∂R must be $+1$ (as opposed to -1). In other words,

$$\partial R = \gamma_{i_1} + \cdots + \gamma_{i_p}$$

for some $1 \leqslant i_1 < \cdots < i_p \leqslant k$, as desired. □

4.3. Prerequisites for contractibility

As we said at the beginning of this chapter, our main result will be that the complex $\mathcal{C}_x(\Sigma)$ is contractible. This will be proven in the next section after we discuss some preliminary results. The heart of our proof will be an explicit deformation retraction of $\widehat{\mathcal{C}}_x(\Sigma)$ to a point; we will then deduce that $\mathcal{C}_x(\Sigma)$ is contractible by giving an explicit (and fairly simple) deformation retraction of $\widehat{\mathcal{C}}_x(\Sigma)$ to $\mathcal{C}_x(\Sigma)$. To construct a deformation retraction of $\widehat{\mathcal{C}}_x(\Sigma)$ to a point, we will construct canonical "straight lines" between any two points in $\widehat{\mathcal{C}}_x(\Sigma)$. This will be done via a parameterization of $\widehat{\mathcal{C}}_x(\Sigma)$ by a set of differential forms; see the map Λ constructed below.

While the entirety of this set of differential forms is not convex, it is close enough to being convex that we can use it to get the desired “straight lines”.

Hyperbolic geometry. This proof is the one place in this book where we will use a tiny amount of hyperbolic geometry. Recall that a hyperbolic metric is a Riemannian metric with constant sectional curvature -1. These exist on all closed surfaces whose genus is at least 2. Fixing a hyperbolic metric on Σ, the following three facts then hold.

- Every nonnullhomotopic simple closed curve on Σ is homotopic to a unique simple geodesic.
- Any two distinct simple geodesics on Σ intersect transversely.
- Let γ and γ' be disjoint nonnullhomotopic simple closed curves on Σ. Assume that γ and γ' are not homotopic to each other. Then the geodesics that are homotopic to γ and γ' are disjoint.

See [**FM12**, §1.1] for this and much more.

Cleaning up curves. If $\gamma_1 \cup \cdots \cup \gamma_k$ is a collection of disjoint oriented simple closed curves on Σ and $c_1, \ldots, c_k \in \mathbb{R}$, then $c_1\gamma_1 + \cdots + c_k\gamma_k$ is not necessarily a positively weighted oriented multicurve : some of the c_i might be negative, some of the γ_i might be homotopic to each other (possibly with opposite orientations), and some of the γ_i might be nullhomotopic. However, by discarding the nullhomotopic γ_i, reversing the orientations of some of the γ_i (and changing the signs of the corresponding c_i), and collecting together the homotopic γ_i, we obtain a canonical positively weighted oriented multicurve c. We will say that c is obtained by *cleaning up* $c_1\gamma_1 + \cdots + c_k\gamma_k$. This definition extends in an obvious way if some of the γ_i are oriented 1-submanifolds with multiple components.

Maps to circle. Consider a smooth map $f : \Sigma \to S^1$. For any regular value $p \in S^1$ of f, the pullback $f^{-1}(p)$ is an oriented 1-submanifold of Σ. We will say that f *represents* the associated element $[f^{-1}(p)]$ of $\mathrm{H}_1(\Sigma;\mathbb{Z})$; this makes sense since $[f^{-1}(p)] = [f^{-1}(q)]$ for any two regular values $p, q \in S^1$. This latter assertion follows from the fact that if λ is an oriented arc of S^1 with oriented boundary $p - q$, then $f^{-1}(\lambda)$ is a subsurface whose oriented boundary is $f^{-1}(p) \sqcup -f^{-1}(q)$. Another way of describing $[f^{-1}(p)]$ is that it is the element of $\mathrm{H}_1(\Sigma;\mathbb{Z})$ which is Poincaré dual to $f^*([S^1]) \in \mathrm{H}^1(\Sigma;\mathbb{Z})$. This can be derived from the relationship between cup products on cohomology and intersection products on homology; see, e.g., [**Bre97**, §VI.12].

Weighted multicurve from map to circle. Now assume that $f : \Sigma \to S^1$ represents $x \in \mathrm{H}_1(\Sigma;\mathbb{Z})$ and has finitely many critical values. These critical values divide S^1 into arcs $\lambda_1, \ldots, \lambda_n$. Normalize S^1 such that its circumference is 1. Setting $c_i = \mathrm{length}(\lambda_i)$, we thus have $c_1 + \cdots + c_n = 1$. For all $1 \leqslant i \leqslant n$, let q_i be an arbitrary point in the interior of λ_i and let $\delta_i = f^{-1}(q_i)$. Since $[\delta_i] = x$ for all $1 \leqslant i \leqslant n$, we have

$$c_1[\delta_1] + \cdots + c_q[\delta_n] = (c_1 + \cdots + c_n)[x] = [x].$$

Define $\Psi(f) \in \widehat{\mathcal{C}}_x(\Sigma)$ to be the result of cleaning up $c_1\delta_1 + \cdots + c_n\delta_n$. The element $\Psi(f)$ appears to depend on the choice of the q_i; however, different choices of q_i will yield homotopic δ_i, and thus $\Psi(f)$ is well-defined.

Globalizing the construction. Define $\mathcal{F}_x(\Sigma, S^1)$ to be the space of smooth maps $\Sigma \to S^1$ representing x which have finitely many critical values. Give $\mathcal{F}_x(\Sigma, S^1)$ the C^∞-topology. The above construction yields a map $\Psi : \mathcal{F}_x(\Sigma, S^1) \to \widehat{\mathcal{C}}_x(\Sigma)$.

LEMMA 4.5. *The map* $\Psi : \mathcal{F}_x(\Sigma, S^1) \to \widehat{\mathcal{C}}_x(\Sigma)$ *is continuous.*

PROOF. As f moves around $\mathcal{F}_x(\Sigma, S^1)$, the critical values of f move continuously around S^1. The 1-submanifolds of Σ used to define $\Psi(f)$ therefore also move homotopically around in Σ. When two critical values come together (causing one of the arcs used to define $\Psi(f)$ to disappear), the weight on the corresponding submanifold of Σ shrinks to 0. □

The following two lemmas show that Ψ is insensitive to certain deformations of its input.

LEMMA 4.6. *Let* $r : S^1 \to S^1$ *be a rotation. Then for all* $f \in \mathcal{F}_x(\Sigma, S^1)$ *we have* $\Psi(f) = \Psi(r \circ f)$.

PROOF. Obvious. □

LEMMA 4.7. *Let* $f_t \in \mathcal{F}_x(\Sigma, S^1)$ *be a continuous family of maps for* $t \in [0, 1]$. *Assume that the critical values of* f_t *and* $f_{t'}$ *are equal for all* $t, t' \in [0, 1]$. *Then* $\Psi(f_0) = \Psi(f_1)$.

PROOF. Let $\lambda_1, \ldots, \lambda_n$ be the arcs into which S^1 is divided by the common critical values of the f_t, and set $c_i = \text{length}(\lambda_i)$. For $1 \leqslant i \leqslant n$, let q_i be an arbitrary point in the interior of λ_i. Finally, for $0 \leqslant t \leqslant 1$ let $\gamma_i(t) = f_t^{-1}(q_i)$. The key observation is that the curve $\gamma_i(t)$ depends continuously on t, so $\gamma_i(0)$ is homotopic to $\gamma_i(1)$. Thus

$$c_1\gamma_1(0) + \cdots + c_n\gamma_n(0) = c_1\gamma_1(1) + \cdots + c_n\gamma_n(1),$$

and the lemma follows. □

Constructing maps using one-forms. To make the above results useful, we need a way of constructing elements of $\mathcal{F}_x(\Sigma, S^1)$. Consider a smooth closed 1-form ω on Σ which is Poincaré dual to $x \in \mathrm{H}_1(\Sigma; \mathbb{Z})$. In other words,

$$\int_h \omega = \hat{i}(x, h) \qquad (h \in \mathrm{H}_1(\Sigma; \mathbb{Z})).$$

For all basepoints $p_0 \in \Sigma$, we can define a smooth map $\Phi(\omega, p_0) : \Sigma \to S^1$ as follows. Regard S^1 as $\mathbb{R}/\mathbb{Z}$. The for any $q \in \Sigma$, we define $\Phi(\omega, p_0)(q)$ to be the image of $\int_\alpha \omega$ in S^1, where α is a smooth path on Σ from p_0 to q. The fact that ω is Poincaré dual to x implies that the integral of ω around any closed loop is an integer, so this is well-defined. The critical points of $\Phi(\omega, p_0)$ are exactly the zeros of ω.

Define $\Omega_x(\Sigma)$ to be the set of smooth closed 1-forms ω on Σ with the following two properties.

- ω is Poincaré dual to $x \in \mathrm{H}_1(\Sigma; \mathbb{Z})$.
- The zero set of ω has finitely many connected components.

Endow $\Omega_x(\Sigma)$ with the C^∞-topology. The above discussion is summarized in the following lemma.

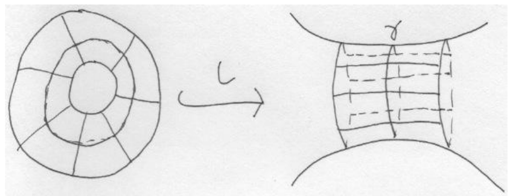

FIGURE 4.5. An ϵ-strip.

LEMMA 4.8. *There exists a continuous map* $\Phi : \Omega_x(\Sigma) \times \Sigma \to \mathcal{F}_x(\Sigma, S^1)$.

PROOF. The only new assertion here is the continuity of Φ, but this is obvious from its definition. □

Changing the basepoint p_0 has the following effect on $\Phi(\omega, p_0)$.

LEMMA 4.9. *Let* $\omega \in \Omega_x(\Sigma)$ *and* $p_0, p_0' \in \Sigma$. *Then* $\Phi(\omega, p_0') = r \circ \Phi(\omega, p_0)$, *where* $r : S^1 \to S^1$ *is a rotation.*

PROOF. We can take r to be a rotation of S^1 by $\int_\alpha \omega$, where α is a smooth path on Σ from p_0' to p_0. □

Combining the constructions. Define a map $\Lambda : \Omega_x(\Sigma) \to \widehat{\mathcal{C}}_x(\Sigma)$ by setting $\Lambda(\omega) = \Psi(\Phi(\omega, p_0))$, where $p_0 \in \Sigma$ is an arbitrary base point. Lemmas 4.6 and 4.9 show that $\Lambda(\omega)$ does not depend on the choice of p_0. The main properties of Λ are contained in the following two lemmas.

LEMMA 4.10. *The map* $\Lambda : \Omega_x(\Sigma) \to \widehat{\mathcal{C}}_x(\Sigma)$ *is continuous.*

PROOF. An immediate consequence of Lemmas 4.5 and 4.8. □

LEMMA 4.11. *For* $0 \leqslant t \leqslant 1$, *let* $\omega_t \in \Omega_x(\Sigma)$ *be a continuous family of* 1-*forms. Assume that we can find a set* $\{p_0, \ldots, p_\ell\}$ *of points on* Σ *with the following properties.*

- *For all* $0 \leqslant t \leqslant 1$, *the set* $\{p_0, \ldots, p_\ell\}$ *consists of exactly one point in each connected component of the zero set of* ω_t.
- *For all* $1 \leqslant i \leqslant \ell$, *there exists an arc* α_i *in* Σ *connecting* p_0 *to* p_i *such that* $\int_{\alpha_i} \omega_t = \int_{\alpha_i} \omega_{t'}$ *for all* $0 \leqslant t, t' \leqslant 1$.

Then $\Lambda(\omega_0) = \Lambda(\omega_1)$.

PROOF. By construction, for all $0 \leqslant t \leqslant 1$ the critical values of $\Phi(\omega_t, p_0)$ are exactly the images in $S^1 = \mathbb{R}/\mathbb{Z}$ of the set $\{0, \int_{\alpha_1} \omega_t, \ldots, \int_{\alpha_\ell} \omega_t\}$. The lemma thus follows from Lemma 4.7. □

Example I : single curve. We now give the first of two examples of the above techniques. Let γ be an oriented simple closed curve on Σ such that $[\gamma] = x$. We will construct some $\omega \in \Omega_x(\Sigma)$ such that $\Lambda(\omega) = \gamma$. Assume that we have fixed a hyperbolic metric on Σ. Homotoping γ, we can assume that it is a geodesic. Parameterize the annulus $\mathbb{A}$ in polar coordinates as $\{(r, \theta) \mid 1 \leqslant r \leqslant 3 \text{ and } 0 \leqslant \theta < 2\pi\}$. For $\epsilon > 0$, an *ϵ-strip map* around γ is an embedding $\iota : \mathbb{A} \hookrightarrow \Sigma$ with the following properties.

- The map ι takes the oriented "core" curve $\{(r,\theta) \mid r = 2,\ 0 \leqslant \theta < 2\pi\}$ of $\mathbb{A}$ to γ, parameterized at constant speed.
- For all angles $0 \leqslant \theta_0 < 2\pi$, the map ι takes the oriented line segment $\{(r,\theta_0) \mid 1 \leqslant r \leqslant 3\}$ in $\mathbb{A}$ to a geodesic segment of length 2ϵ that intersects γ orthogonally with a positive sign. Again, this geodesic segment is parameterized at constant speed.

See Figure 4.5. For $\epsilon > 0$ sufficiently small these exist and are unique up to precomposition with a rotation of $\mathbb{A}$. The image A of ι will be called an *ϵ-strip* around γ. Define $\mu : \mathbb{R} \to \mathbb{R}$ to be the function

$$\mu(x) = \begin{cases} \frac{1}{\int_{-\infty}^{\infty} e^{-1/(1-(z-2)^2)}\mathrm{d}z} \int_0^x e^{-1/(1-(z-2)^2)}\mathrm{d}z & \text{if } x \in [1,3], \\ 0 & \text{if } x \notin [1,3]. \end{cases}$$

Thus μ is a smooth nonnegative function of total integral 1 which is supported on $[1,3]$. There is a smooth closed 1-form $\mu(r)\mathrm{d}r$ on $\mathbb{A}$. We can therefore define a smooth closed 1-form ω on Σ via the formulas

$$\omega|_A = \iota_*(\mu(r)\mathrm{d}r) \quad \text{and} \quad \omega|_{\Sigma\setminus A} = 0.$$

We will call ω the *ϵ-strip form* dual to γ. It is clear that ω represents $[\gamma] = x$. Additionally, we have the following lemma

LEMMA 4.12. *With the notation as above, we have $\Lambda(\omega) = \gamma$.*

PROOF. Fix a basepoint $p_0 \in \Sigma\setminus A$. Regarding S^1 as $\mathbb{R}/\mathbb{Z}$, it is then clear from the definitions that $\Phi(\omega, p_0) : \Sigma \to S^1$ is the map

$$\Phi(\omega,p_0)(q) = \begin{cases} \int_1^r \mu(r)\mathrm{d}r & \text{if } q = \iota(r,\theta) \in A \text{ with } (r,\theta) \in \mathbb{A}, \\ 0 & \text{if } q \notin A. \end{cases}$$

In particular, the only critical value of $\Phi(\omega,p_0)$ is 0, and the preimage under $\Phi(\omega,p_0)$ of a regular value $q_1 \in (0,1) \subset S^1$ is a loop of the form $\{\iota(r_1,\theta) \mid 0 \leqslant \theta < 2\pi\}$ for some $1 < r_1 < 3$. This loop is homotopic to γ, so we conclude that

$$\Lambda(\omega) = \Psi(\Phi(\omega,p_0)) = 1 \cdot \lambda = \lambda,$$

as desired. □

Example II : multicurve. We now generalize the previous example. Let $c = c_1\gamma_1 + \cdots + c_k\gamma_k$ be an arbitrary positively weighted oriented multicurve on Σ which represents x. Again assume that we have fixed a hyperbolic metric on Σ and that each γ_i is a geodesic. Let $\epsilon > 0$ be small enough that there are ϵ-strips around each γ_i which are pairwise disjoint. For $1 \leqslant i \leqslant k$ let ω_i be the ϵ-strip form dual to γ_i, so ω_i represents $[\gamma_i]$. Finally, define $\omega = c_1\omega_1 + \cdots + c_k\omega_k$. It is then an easy exercise in the definitions to see that ω represents $c_1[\gamma_1] + \cdots + c_k[\gamma_k]$. Moreover, we have the following generalization of Lemma 4.12.

LEMMA 4.13. *With the notation as above, we have $\Lambda(\omega) = c$.*

PROOF. For $1 \leqslant i \leqslant k$, let A_i be the ϵ-strip around γ_i. Pick a basepoint $p_0 \in \Sigma\setminus \cup_{i=1}^k A_i$. Just like in the proof of Lemma 4.12, the map $\Phi(\Sigma,p_0) : \Sigma \to S^1$ takes p_0 to $0 \in S^1 = \mathbb{R}/\mathbb{Z}$, takes each component of $\Sigma\setminus \cup_{i=1}^k A_i$ to a critical value, and takes A_i to an arc of $S^1 = \mathbb{R}/\mathbb{Z}$ of length c_i (starting and ending at a critical value; observe that this arc can contain critical values in its interior). Let $\lambda_1, \ldots, \lambda_n$

be the arcs into which S^1 is divided by the critical values, let $d_i = \text{length}(\lambda_i)$, and let q_i be an arbitrary point in the interior of λ_i. Define $\delta_i = \Phi(\Sigma, p_0)^{-1}(q_i)$, so $\Lambda(\omega)$ is the result of cleaning up

$$d_1\delta_1 + \cdots + d_n\delta_n. \tag{6}$$

It is clear that $\cup_{j=1}^n \delta_j \subset \cup_{i=1}^k A_i$.

Fix some $1 \leqslant i \leqslant k$. The components of $\delta_1 \cup \cdots \cup \delta_n$ lying in A_i consist of a set of curves each of which is homotopic to γ_i. From (6), each of these curves has a weight from among the numbers $d_1, \ldots, d_n$. It is easy to see that these weights add up to c_i. The lemma follows. □

4.4. Contractibility

We finally prove the following theorem of Bestvina–Bux–Margalit [**BBM10**]. In its proof, we will use all of the notation introduced in §4.3.

THEOREM 4.14. *Let Σ be a closed surface and let $x \in \mathrm{H}_1(\Sigma;\mathbb{Z})$ be a primitive vector. Then $\mathcal{C}_x(\Sigma)$ is contractible.*

PROOF. The theorem has no content if the genus of Σ is 0 since in that case $\mathrm{H}_1(\Sigma;\mathbb{Z}) = 0$ contains no primitive vectors. If the genus of Σ is 1, then Σ contains no oriented multicurves with more than one component. The complex $\mathcal{C}_x(\Sigma)$ therefore is a discrete set of points, one for each homotopy class of oriented simple closed curve γ with $[\gamma] = x$. Theorem 2.8 says that there is at least one such curve, and Proposition 3.24 says that the Torelli group $\mathcal{I}(\Sigma)$ acts transitively on them. However, Lemma 3.2 says that $\mathcal{I}(\Sigma) = 1$. We conclude that if the genus of Σ is 1, then $\mathcal{C}_x(\Sigma)$ consists of exactly one point and is hence contractible.

We can therefore assume without loss of generality that Σ has genus at least 2, which allows us to fix a hyperbolic metric on Σ. The proof of the theorem now has two steps.

STEP 1. *The space $\widehat{\mathcal{C}}_x(\Sigma)$ is contractible.*

Using Theorem 2.8, we can find an oriented simple closed curve γ_0 on Σ such that $[\gamma_0] = x$. Homotoping γ_0, we can assume that it is a hyperbolic geodesic. We will construct an explicit homotopy $f_t : \widehat{\mathcal{C}}_x(\Sigma) \to \widehat{\mathcal{C}}_x(\Sigma)$ such that $f_0 = \text{id}$ and such that $f_1(c) = \gamma_0$ for all $c \in \widehat{\mathcal{C}}_x(\Sigma)$. This construction is divided into three substeps. In the first, we construct f_t on a fixed cell $\mathcal{X}_x(\gamma)$. This construction depends on a parameter $\epsilon > 0$; the second substep shows that in fact its output is independent of ϵ. The final substep shows how to piece together the maps on the various cells to define f_t.

SUBSTEP 1.1. *Let γ be an oriented multicurve such that $\mathcal{X}_x(\gamma)$ is nondegenerate. For all $\epsilon > 0$ sufficiently small, we construct a homotopy $f^{\epsilon}_{\gamma,t} : \mathcal{X}_x(\gamma) \to \widehat{\mathcal{C}}_x(\Sigma)$ such that $f^{\epsilon}_{\gamma,0}$ is the inclusion and $f^{\epsilon}_{\gamma,1}(c) = \gamma_0$ for all $c \in \mathcal{X}_x(\gamma)$.*

Write $\gamma = \gamma_1 \cup \cdots \cup \gamma_k$. Homotoping the γ_i, we can assume that they are all hyperbolic geodesics. For $0 \leqslant i \leqslant k$, let A_i^{ϵ} be an ϵ-strip around γ_i. Choosing $\epsilon > 0$ small enough, we can assume that that the following hold.

- For $1 \leqslant i < j \leqslant k$, we have $A_i^{\epsilon} \cap A_j^{\epsilon} = \varnothing$.
- For $1 \leqslant i \leqslant k$, the ϵ-strips A_0^{ϵ} and A_i^{ϵ} intersect transversely as in Figure 4.6.

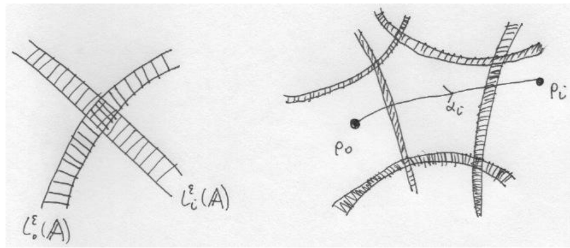

FIGURE 4.6. Left : Two transverse ϵ-strips. Right : The arc α_i crosses some of the ϵ-strips.

For $0 \leqslant i \leqslant k$, let ω_i^ϵ be the ϵ-strip form dual to γ_i. For a point $c_1\gamma_1 + \cdots + c_k\gamma_k$ of $\mathcal{X}_x(\gamma)$ and some $0 \leqslant t \leqslant 1$, the 1-form

$$t\omega_0^\epsilon + (1-t)c_1\omega_1^\epsilon + \cdots + (1-t)c_k\omega_k^\epsilon \tag{7}$$

represents x. Moreover, the following hold.

- For $t = 0$, the zero set of (7) is the complement of $A_1^\epsilon \cup \cdots \cup A_k^\epsilon$.
- For $0 < t < 1$, the zero set of (7) is the complement of $A_0^\epsilon \cup \cdots \cup A_k^\epsilon$ (this follows from our assumptions on the intersections of A_0^ϵ and A_i^ϵ for $1 \leqslant i \leqslant k$).
- For $t = 1$, the zero set of (7) is the complement of A_0^ϵ.

In particular, the zero set of (7) has finitely many components. The upshot of all of this is that (7) is an element of $\Omega_x(\Sigma)$ for all $0 \leqslant t \leqslant 1$. We can therefore define a function $f_{\gamma,t}^\epsilon : \mathcal{X}_x(\gamma) \to \widehat{\mathcal{C}}_x(\Sigma)$ via the formula

$$f_{\gamma,t}^\epsilon(c_1\gamma_1 + \cdots + c_k\gamma_k) = \Lambda(t\omega_0^\epsilon + (1-t)c_1\omega_1^\epsilon + \cdots + (1-t)c_k\omega_k^\epsilon).$$

Lemma 4.10 implies that $f_{\gamma,t}^\epsilon$ is continuous (both as a function and as a homotopy). Also, it follows from Lemma 4.13 that $f_{\gamma,0}^\epsilon$ is the inclusion and $f_{\gamma,1}^\epsilon(c) = \gamma_0$ for all $c \in \mathcal{X}_x(\gamma)$.

SUBSTEP 1.2. *Let γ be an oriented multicurve such that $\mathcal{X}_x(\gamma)$ is nondegenerate and let $\epsilon, \epsilon' > 0$ be small enough that $f_{\gamma,t}^\epsilon$ and $f_{\gamma,t}^{\epsilon'}$ are defined. Then $f_{\gamma,t}^\epsilon = f_{\gamma,t}^{\epsilon'}$.*

Without loss of generality, $\epsilon' < \epsilon$. As in Substep 1.1, write $\gamma = \gamma_1 \cup \cdots \cup \gamma_k$ with γ_i a hyperbolic geodesic for $1 \leqslant i \leqslant k$. Fix some $c_1\gamma_1 + \cdots + c_k\gamma_k \in \mathcal{X}_x(\gamma)$ and some $0 \leqslant t_0 \leqslant 1$. Our goal is to show that

$$f_{\gamma,t_0}^\epsilon(c_1\gamma_1 + \cdots + c_k\gamma_k) = f_{\gamma,t_0}^{\epsilon'}(c_1\gamma_1 + \cdots + c_k\gamma_k). \tag{8}$$

To simplify our notation, we will deal with the case where $0 < t_0 < 1$; the other cases are similar.

We now set up some notation. For $0 \leqslant i \leqslant k$ let A_i^ϵ be an ϵ-strip around γ_i. Also, for $0 \leqslant i \leqslant k$ and $\epsilon' \leqslant e \leqslant \epsilon$ let ω_i^e be the e-strip form dual to γ_i. Finally, for $\epsilon' \leqslant e \leqslant \epsilon$ let

$$\omega^e = t\omega_0^e + (1-t)c_1\omega_1^e + \cdots + (1-t)c_k\omega_k^e.$$

The assertion of (8) is thus equivalent to the assertion that $\Lambda(\omega^\epsilon) = \Lambda(\omega^{\epsilon'})$.

We will prove this using Lemma 4.11, whose conditions we now verify. First, by construction the differential forms ω^e depend continuously on e. Let $\{p_0, \ldots, p_\ell\}$

be a set of points on Σ that contains exactly one point in the interior of each component of

$$\Sigma \setminus \bigcup_{i=0}^{k} A_i^{\epsilon}.$$

Clearly $\{p_0, \ldots, p_\ell\}$ also contains exactly one point in the interior of each component of

$$\Sigma \setminus \bigcup_{i=0}^{k} A_i^{e}$$

for each $\epsilon' \leqslant e \leqslant \epsilon$. As we said in Substep 1.1, these are exactly the components of the zero set of ω^e (this is where we use the fact that $0 < t_0 < 1$). Finally, for $1 \leqslant i \leqslant \ell$ let α_i be any smooth arc from p_0 to p_i that crosses the γ_i transversely. Letting $\hat{i}(\alpha_i, \gamma_j)$ be the algebraic intersection number between the arc α_i and the simple closed curve γ_j, it is clear that for $1 \leqslant i \leqslant \ell$ and $\epsilon' \leqslant e \leqslant \epsilon$ we have

$$\int_{\alpha_i} \omega^e = t\hat{i}(\alpha_i, \gamma_0) + \sum_{j=1}^{k} (1-t) c_i \hat{i}(\alpha_i, \gamma_j).$$

See Figure 4.6. As this does not depend on e, the conditions of Lemma 4.11 are satisfied and we conclude that $\Lambda(\omega^\epsilon) = \Lambda(\omega^{\epsilon'})$, as desired.

Substep 1.3. *We construct a homotopy $f_t : \widehat{\mathcal{C}}_x(\Sigma) \to \widehat{\mathcal{C}}_x(\Sigma)$ such that $f_0 = id$ and such that $f_1(c) = \delta_0$ for all $c \in \widehat{\mathcal{C}}_x(\Sigma)$.*

If γ is any oriented multicurve such that $\mathcal{X}_x(\gamma)$ is nondegenerate, then using Substep 1.2 we can write $f_{\gamma,t} : \mathcal{X}_x(\gamma) \to \widehat{\mathcal{C}}_x(\Sigma)$ for $f_{\gamma,t}^{\epsilon}$, where $\epsilon > 0$ is an sufficiently small number. To show that the $f_{\gamma,t}$ piece together to give a function $f_t : \widehat{\mathcal{C}}_x(\Sigma) \to \widehat{\mathcal{C}}_x(\Sigma)$, it is enough to show that if γ and γ' are any oriented multicurves such that $\mathcal{X}_x(\gamma)$ and $\mathcal{X}_x(\gamma')$ are nondegenerate, then $f_{\gamma,t}$ and $f_{\gamma',t}$ agree on the intersection of $\mathcal{X}_x(\gamma)$ and $\mathcal{X}_x(\gamma')$ in $\widehat{\mathcal{C}}_x(\Sigma)$. If this intersection is nonempty, then it is exactly $\mathcal{X}_x(\gamma'')$, where γ'' is the oriented multicurve consisting of all oriented simple closed curves that appear in both γ and γ'. But it is clear from their definitions that if $\epsilon > 0$ is small enough that all three of $f_{\gamma,t}^{\epsilon}$ and $f_{\gamma',t}^{\epsilon}$ and $f_{\gamma'',t}^{\epsilon}$ are defined, then all three of them agree on $\mathcal{X}_x(\gamma'')$.

Step 2. *The space $\widehat{\mathcal{C}}_x(\Sigma)$ deformation retracts to $\mathcal{C}_x(\Sigma) \subset \widehat{\mathcal{C}}_x(\Sigma)$.*

Consider a point $c = c_1\gamma_1 + \cdots + c_k\gamma_k$ in $\widehat{\mathcal{C}}_x(\Sigma)$. Discarding some the the γ_i, we can assume that $c_i > 0$ for all $1 \leqslant i \leqslant k$. We will write down a canonical (i.e. independent of all choices) path from c to $\mathcal{C}_x(\Sigma)$. It will be clear that this path depends continuously on c and that it is the constant path if $c \in \mathcal{C}_x(\Sigma)$.

If $c \notin \mathcal{C}_x(\Sigma)$, then the cell $\mathcal{X}_x(\gamma)$ is not reduced. Lemma 4.4 therefore implies that there exists some subsurface R of Σ such that

$$\partial R = \gamma_{i_1} + \cdots + \gamma_{i_p}$$

for some $1 \leqslant i_1 < \cdots < i_p \leqslant k$. Let $R_1, \ldots, R_q$ be *all* such subsurfaces. It follows that

$$\partial R_1 + \cdots + \partial R_q = d_1\gamma_1 + \cdots + d_k\gamma_k$$

for some $d_i \geqslant 0$ (not all 0). Setting $T = \min\{c_i/d_i \mid d_i > 0\}$, we have a path

$$t \mapsto c - t(\partial R_1 + \cdots + \partial R_q)$$

in $\widehat{\mathcal{C}}_x(\Sigma)$ defined for $0 \leqslant t \leqslant T$. At the endpoint of this path, the coefficient of at least one of the γ_i has become 0. Repeat this process until c ends up in $\mathcal{C}_x(\Sigma)$. $\square$

CHAPTER 5

Stabilizers of simple closed curves

Let γ be an oriented nonseparating simple closed curve on Σ_g. This chapter is devoted to the stabilizer $(\mathcal{I}_g)_\gamma$ of γ in $\mathcal{I}_g$. Most of the hard work here is devoted to understanding the symplectic representation of the stabilizer $(\text{Mod}_g)_\gamma$ of γ in the mapping class group. This lands in the stabilizer in $\text{Sp}_{2g}(\mathbb{Z})$ of the homology class $[\gamma] \in \text{H}_1(\Sigma_g; \mathbb{Z})$, and §5.1 is devoted to understanding this stabilizer subgroup. Next, in §5.2 we discuss the mapping class group, and finally in §5.3 we discuss the Torelli group. The results in this chapter are due to van den Berg [**vdB03**] and Putman [**Put07**].

5.1. Stabilizers in the symplectic group

Pick a symplectic basis $\{a_1, b_1, \ldots, a_g, b_g\}$ for $\text{H}_1(\Sigma_g; \mathbb{Z})$. Using this basis, we will identify $\text{H}_1(\Sigma_g; \mathbb{Z})$ with $\mathbb{Z}^{2g}$. The goal of this section is to understand the stabilizer subgroup $(\text{Sp}_{2g}(\mathbb{Z}))_{b_g}$ of b_g.

Special linear group. As motivation, we begin by discussing an analogous stabilizer subgroup in $\text{SL}_n(\mathbb{Z})$. Let $\{\vec{e}_1, \ldots, \vec{e}_n\}$ be the standard basis for $\mathbb{Z}^n$. We then have

$$(\text{SL}_n(\mathbb{Z}))_{\vec{e}_n} = \{ \left(\begin{array}{ccc|c} & & & 0 \\ & A & & \vdots \\ & & & 0 \\ \hline c_1 & \cdots & c_{n-1} & 1 \end{array}\right) \mid A \in \text{SL}_{n-1}(\mathbb{Z}),\ c_1, \ldots, c_{n-1} \in \mathbb{Z}\}.$$

This decomposes as

$$(\text{SL}_n(\mathbb{Z}))_{\vec{e}_n} = \mathbb{Z}^{n-1} \rtimes \text{SL}_{n-1}(\mathbb{Z}), \tag{9}$$

where $\text{SL}_{n-1}(\mathbb{Z})$ is embedded in $\text{SL}_n(\mathbb{Z})$ in the usual way and

$$\mathbb{Z}^{n-1} = \{ \left(\begin{array}{ccc|c} & & & 0 \\ & I & & \vdots \\ & & & 0 \\ \hline c_1 & \cdots & c_{n-1} & 1 \end{array}\right) \mid c_1, \ldots, c_{n-1} \in \mathbb{Z}\}.$$

The action of $\text{SL}_{n-1}(\mathbb{Z})$ on the abelian group $\mathbb{Z}^{n-1}$ is the obvious one. The associated projection $\rho : (\text{SL}_n(\mathbb{Z}))_{\vec{e}_n} \to \text{SL}_{n-1}(\mathbb{Z})$ can be described as follows. Consider some $\phi \in (\text{SL}_n(\mathbb{Z}))_{\vec{e}_n}$. For $v \in \mathbb{Z}^{n-1}$, we have $\phi(v) = v' + p\vec{e}_n$ for some $v' \in \mathbb{Z}^{n-1}$ and $p \in \mathbb{Z}$. We then have that $\rho(\phi)(v) = v'$. From this point of view, we see that the kernel $\mathbb{Z}^{n-1}$ of ρ is exactly the subgroup of all $\phi \in (\text{SL}_n(\mathbb{Z}))_{\vec{e}_n}$ such that for all $v \in \mathbb{Z}^{n-1}$, we have $\phi(v) = v + p\vec{e}_n$ for some $p \in \mathbb{Z}$.

Symplectic group. We now return to the symplectic group. We begin with the following lemma. Identify $\mathbb{Z}^{2(g-1)}$ with the subgroup $\langle a_1, b_1, \ldots, a_{g-1}, b_{g-1}\rangle$ of $\mathbb{Z}^{2g}$.

LEMMA 5.1. *Consider $\phi \in (\mathrm{Sp}_{2g}(\mathbb{Z}))_{b_g}$. Then for all $v \in \mathbb{Z}^{2(g-1)}$, we have $\phi(v) = v' + pb_g$ for some $v' \in \mathbb{Z}^{2(g-1)}$ and $p \in \mathbb{Z}$.*

PROOF. Write $\phi(v) = v' + p'a_g + pb_g$, where $v' \in \mathbb{Z}^{2(g-1)}$ and $p, p' \in \mathbb{Z}$. We then have

$$0 = \hat{i}(v, b_g) = \hat{i}(\phi(v), \phi(b_g)) = \hat{i}(v' + p'a_g + pb_g, b_g) = p',$$

as desired. □

We now define a homomorphism $\pi : (\mathrm{Sp}_{2g}(\mathbb{Z}))_{b_g} \to \mathrm{Sp}_{2(g-1)}(\mathbb{Z})$ as follows. Consider $\phi \in (\mathrm{Sp}_{2g}(\mathbb{Z}))_{b_g}$ and $v \in \mathbb{Z}^{2(g-1)}$. Using Lemma 5.1, write $\phi(v) = v' + pb_g$ with $v' \in \mathbb{Z}^{2(g-1)}$ and $p \in \mathbb{Z}$. We then define $\pi(\phi)(v) = v'$. It is easy to see that π is a homomorphism whose image lies in $\mathrm{Sp}_{2(g-1)}(\mathbb{Z})$. Define $K_g = \ker(\pi)$. Thus K_g is the subgroup of $(\mathrm{Sp}_{2g}(\mathbb{Z}))_{b_g}$ consisting of all $\phi \in (\mathrm{Sp}_{2g}(\mathbb{Z}))_{b_g}$ such that for all $v \in \mathbb{Z}^{2(g-1)}$, we have $\phi(v) = v + pb_g$ for some $p \in \mathbb{Z}$. We have a short exact sequence

$$1 \longrightarrow K_g \longrightarrow (\mathrm{Sp}_{2g}(\mathbb{Z}))_{b_g} \xrightarrow{\pi} \mathrm{Sp}_{2(g-1)}(\mathbb{Z}) \longrightarrow 1$$

which splits via the standard inclusion $\mathrm{Sp}_{2(g-1)}(\mathbb{Z}) \hookrightarrow (\mathrm{Sp}_{2g}(\mathbb{Z}))_{b_g}$. We thus have proved the following proposition.

PROPOSITION 5.2. *If $\{a_1, b_1, \ldots, a_g, b_g\}$ is a symplectic basis for $\mathrm{H}_1(\Sigma_g; \mathbb{Z})$, then we have $(\mathrm{Sp}_{2g}(\mathbb{Z}))_{b_g} = K_g \rtimes \mathrm{Sp}_{2(g-1)}(\mathbb{Z})$.*

The kernel. Unlike in $\mathrm{SL}_n(\mathbb{Z})$, the kernel group K_g is *not* abelian. Rather, we will soon see that it is 2-step nilpotent. The key to this is the following lemma.

LEMMA 5.3. *Consider $\phi \in K_g$. We then have $\phi(a_g) = w + a_g + qb_g$ for some $w \in \mathbb{Z}^{2(g-1)}$ and $q \in \mathbb{Z}$. Moreover, for $v \in \mathbb{Z}^{2(g-1)}$ we have $\phi(v) = v + \hat{i}(v, w)b_g$.*

PROOF. We can write $\phi(a_g) = w + q'a_g + qb_g$ for some $w \in \mathbb{Z}^{2(g-1)}$ and $q, q' \in \mathbb{Z}$. Then

$$1 = \hat{i}(a_g, b_g) = \hat{i}(\phi(a_g), \phi(b_g)) = \hat{i}(w + q'a_g + qb_g, b_g) = q',$$

as desired. For the second part of the lemma, we know by the definition of K_g that for $v \in \mathbb{Z}^{2(g-1)}$ we have $\phi(v) = v + pb_g$ for some $p \in \mathbb{Z}$. We then have

$$0 = \hat{i}(v, a_g) = \hat{i}(\phi(v), \phi(a_g)) = \hat{i}(v + pb_g, w + a_g + qb_g) = \hat{i}(v, w) - p,$$

so $p = \hat{i}(v, w)$. □

For $w \in \mathbb{Z}^{2(g-1)}$ and $q \in \mathbb{Z}$, define $\phi_{w,q} : \mathbb{Z}^{2g} \to \mathbb{Z}^{2g}$ via the formulas

$$\phi_{w,q}(a_g) = w + a_g + qb_g \quad \text{and} \quad \phi_{w,q}(b_g) = b_g$$

and

$$\phi_{w,q}(v) = v + \hat{i}(v, w)b_g \qquad (v \in \mathbb{Z}^{2(g-1)}).$$

It is clear that $\phi_{w,q} \in K_g$, and by Lemma 5.3 every element of K_g is of the form $\phi_{w,q}$ for some unique $w \in \mathbb{Z}^{2(g-1)}$ and $q \in \mathbb{Z}$. We now make the following calculation.

LEMMA 5.4. *For $w_1, w_2 \in \mathbb{Z}^{2(g-1)}$ and $q_1, q_2 \in \mathbb{Z}$, we have*

$$\phi_{w_1,q_1}\phi_{w_2,q_2} = \phi_{w_1+w_2,q_1+q_2+\hat{i}(w_2,w_1)}.$$

PROOF. Using the definition of the ϕ_{w_i,q_i}, we calculate as follows.

$$\begin{aligned}\phi_{w_1,q_1}(\phi_{w_2,q_2}(a_g)) &= \phi_{w_1,q_1}(w_2 + a_g + q_2 b_g)\\ &= (w_2 + \hat{i}(w_2,w_1)b_g) + (w_1 + a_g + q_1 b_g) + q_2 b_g\\ &= (w_1 + w_2) + a_g + (q_1 + q_2 + \hat{i}(w_2,w_1))b_g.\end{aligned}$$

The lemma follows. □

We can now prove the following.

PROPOSITION 5.5. *There is a nonsplit central extension*

$$1 \longrightarrow \mathbb{Z} \longrightarrow K_g \longrightarrow \mathbb{Z}^{2(g-1)} \longrightarrow 1,$$

where the central $\mathbb{Z}$ consists of $\{\phi_{0,q} \mid q \in \mathbb{Z}\}$.

PROOF. Using Lemma 5.4, we can define a homomorphism $\pi : K_g \to \mathbb{Z}^{2(g-1)}$ via the formula $\pi(\phi_{w,q}) = w$. The kernel of π is exactly $\{\phi_{0,q} \mid q \in \mathbb{Z}\}$, which is in the center of K_g by Lemma 5.4. To see that the resulting central extension is not split, it is enough to observe that K_g is not abelian, which is immediate from Lemma 5.4. □

REMARK 5.6. One might thing that a splitting $\mathbb{Z}^{2(g-1)} \to K_g$ can be defined by the formula $w \mapsto \phi_{w,0}$. However, the formulas in Lemma 5.4 show that this is not actually a homomorphism. For instance,

$$\phi_{b_1,0}\phi_{a_1,0} = \phi_{a_1+b_1,1}.$$

5.2. The symplectic representation of $\mathrm{Mod}_{g,\gamma}$

We now turn to the mapping class group. Fix $g \geqslant 2$, and let γ be an oriented nonseparating simple closed curve on Σ_g.

Cut open surface. We begin by recalling the notation and results from §1.4. Recall that $\Sigma_{g,\gamma}$ is the surface obtained by cutting Σ_g along γ and $\mathrm{Mod}_{g,\gamma}$ is the mapping class group of $\Sigma_{g,\gamma}$. Letting $\{\partial_1, \partial_2\}$ be the boundary components of $\Sigma_{g,\gamma}$, Lemma 1.20 says that there is a short exact sequence

$$1 \longrightarrow \mathbb{Z} \longrightarrow \mathrm{Mod}_{g,\gamma} \longrightarrow (\mathrm{Mod}_g)_\gamma \longrightarrow 1, \tag{10}$$

where $\mathbb{Z}$ is generated by $T_{\partial_1}T_{\partial_2}^{-1}$. Recall that a γ-splitting surface in $\Sigma_{g,\gamma}$ is a subsurface S of $\Sigma_{g,\gamma}$ such that $\Sigma_{g,\gamma} \setminus S$ is a 3-holed sphere two of whose boundary components are $\{\partial_1, \partial_2\}$. Fix a γ-splitting surface S. Letting $\widehat{\Sigma}_{g,\gamma}$ be the result of gluing a disc to $\Sigma_{g,\gamma}$ along ∂_1, Lemma 1.21 says that there is a decomposition

$$\mathrm{Mod}_{g,\gamma} = \pi_1(U\widehat{\Sigma}_{g,\gamma}) \rtimes \mathrm{Mod}(S). \tag{11}$$

One should view (11) as being analogous to Proposition 5.2 above.

Symplectic representation. The map $\mathrm{H}_1(\Sigma_{g,\gamma};\mathbb{Z}) \to \mathrm{H}_1(\Sigma_g;\mathbb{Z})$ induced by the map $\Sigma_{g,\gamma} \to \Sigma_g$ that glues ∂_1 and ∂_2 back together is injective, and we will identify $\mathrm{H}_1(\Sigma_{g,\gamma};\mathbb{Z})$ with its image in $\mathrm{H}_1(\Sigma_g;\mathbb{Z})$. Also, the map $\mathrm{H}_1(S;\mathbb{Z}) \to \mathrm{H}_1(\Sigma_{g,\gamma};\mathbb{Z})$ induced by the inclusion $S \hookrightarrow \Sigma_{g,\gamma}$ is injective, and we will identify $\mathrm{H}_1(S;\mathbb{Z})$ with

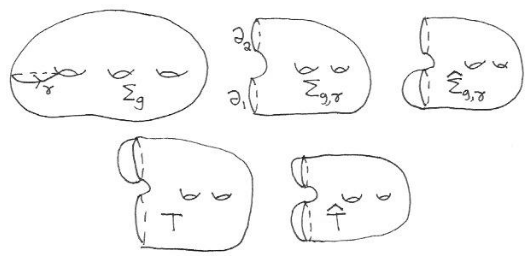

FIGURE 5.1. The various surfaces involved in the proof of Lemma 5.7

its image in $\mathrm{H}_1(\Sigma_{g,\gamma};\mathbb{Z}) \subset \mathrm{H}_1(\Sigma_g;\mathbb{Z})$. Choose a symplectic basis $\{a_1, b_1, \ldots, a_g, b_g\}$ for $\mathrm{H}_1(\Sigma_g;\mathbb{Z})$ such that $b_g = [\gamma]$ and such that $\{a_1, b_1, \ldots, a_{g-1}, b_{g-1}\}$ is a symplectic basis for $\mathrm{H}_1(S;\mathbb{Z}) \subset \mathrm{H}_1(\Sigma_g;\mathbb{Z})$. Composing the surjection $\mathrm{Mod}_{g,\gamma} \to (\mathrm{Mod}_g)_\gamma$ with the symplectic representation $\mathrm{Mod}_g \to \mathrm{Sp}_{2g}(\mathbb{Z})$, we obtain a homomorphism

$$\psi : \mathrm{Mod}_{g,\gamma} \longrightarrow (\mathrm{Sp}_{2g}(\mathbb{Z}))_{b_g}.$$

Semidirect product decompositions and the symplectic representation. Proposition 5.2 says that there is a decomposition

$$(\mathrm{Sp}_{2g}(\mathbb{Z}))_{b_g} = K_g \rtimes \mathrm{Sp}_{2(g-1)}(\mathbb{Z}). \tag{12}$$

Our first result says that ψ "respects" the semidirect product decompositions (11) and (12).

LEMMA 5.7. *We have $\psi(\pi_1(U\widehat{\Sigma}_{g,\gamma})) \subset K_g$ and $\psi(\mathrm{Mod}(S)) \subset \mathrm{Sp}_{2(g-1)}(\mathbb{Z})$.*

PROOF. The fact that $\psi(\mathrm{Mod}(S)) \subset \mathrm{Sp}_{2(g-1)}(\mathbb{Z})$ is an immediate consequence of the fact that the set $\{a_1, b_1, \ldots, a_{g-1}, b_{g-1}\}$ is a symplectic basis for $\mathrm{H}_1(S;\mathbb{Z}) \subset \mathrm{H}_1(\Sigma_g;\mathbb{Z})$.

We now prove that $\psi(\pi_1(U\widehat{\Sigma}_{g,\gamma})) \subset K_g$. To keep the various surfaces constructed in this part of the proof straight, we recommend that the reader consult Figure 5.1. We showed in Lemma 5.1 that $(\mathrm{Sp}_{2g}(\mathbb{Z}))_{b_g}$ preserves the subspace $\mathrm{H}_1(\Sigma_{g,\gamma};\mathbb{Z}) = \langle a_1, b_1, \ldots, a_{g-1}, b_{g-1}, b_g\rangle$. By definition, K_g is the group consisting of elements of $(\mathrm{Sp}_{2g}(\mathbb{Z}))_{b_g}$ that act trivially on the quotient of $\mathrm{H}_1(\Sigma_{g,\gamma};\mathbb{Z})$ by $\langle b_g\rangle = [\partial_2]$. Letting T be the surface obtained by gluing a disc to $\Sigma_{g,\gamma}$ along ∂_2, the map $\mathrm{H}_1(\Sigma_{g,\gamma};\mathbb{Z}) \to \mathrm{H}_1(T;\mathbb{Z})$ induced by the inclusion $T \hookrightarrow \Sigma_{g,\gamma}$ is a surjection whose kernel is spanned by $[\partial_2]$. There is a map $\mathrm{Mod}_{g,\gamma} \to \mathrm{Mod}(T)$ that extends mapping classes over the glued-in disc by the identity, and by what we have said it is enough to show that the image of $\pi_1(U\widehat{\Sigma}_{g,\gamma}) \subset \mathrm{Mod}_{g,\gamma}$ in $\mathrm{Mod}(T)$ acts trivially on $\mathrm{H}_1(T;\mathbb{Z})$ (i.e. lies in $\mathcal{I}(T)$).

Letting $\widehat{T}$ be the result of gluing a disc to T along ∂_1, there is a Birman exact sequence

$$1 \longrightarrow \pi_1(U\widehat{T}) \longrightarrow \mathrm{Mod}(T) \longrightarrow \mathrm{Mod}(\widehat{T}) \longrightarrow 1;$$

see Theorem 1.17. Clearly the image of the disc-pushing subgroup $\pi_1(U\widehat{\Sigma}_{g,\gamma})$ in $\mathrm{Mod}(T)$ lies in the disc-pushing subgroup $\pi_1(U\widehat{T})$, so it is enough to show that $\pi_1(U\widehat{T}) \subset \mathcal{I}(T)$. We showed this in Theorem 3.31; recall that the key point is that the inclusion map $T \hookrightarrow \widehat{T}$ induces an isomorphism $\mathrm{H}_1(T;\mathbb{Z}) \cong \mathrm{H}_1(\widehat{T};\mathbb{Z})$, so the symplectic representation of $\mathrm{Mod}(T)$ factors through $\mathrm{Mod}(\widehat{T})$. The lemma follows. $\square$

Disc-pushing and the symplectic representation. We now investigate the restriction of the symplectic representation $\psi : \mathrm{Mod}_{g,\gamma} \to (\mathrm{Sp}_{2g}(\mathbb{Z}))_{b_g}$ to the disc-pushing subgroup $\pi_1(U\widehat{\Sigma}_{g,\gamma})$, which Lemma 5.7 says lands in K_g. The group $\pi_1(U\widehat{\Sigma}_{g,\gamma})$ fits into an exact sequence

$$1 \longrightarrow \mathbb{Z} \longrightarrow \pi_1(U\widehat{\Sigma}_{g,\gamma}) \longrightarrow \pi_1(\widehat{\Sigma}_{g,\gamma}) \longrightarrow 1, \tag{13}$$

where the kernel $\mathbb{Z}$ is generated by the loop around the fiber. Recall from §1.4 that the loop around the fiber corresponds to the mapping class $T_{\partial_1} \in \mathrm{Mod}_{g,\gamma}$.

The following lemma says that the exact sequence (13) is compatible with the one given by Proposition 5.5. In its statement, we are identifying $\mathrm{H}_1(\widehat{\Sigma}_{g,\gamma})$ with $\mathrm{H}_1(S;\mathbb{Z}) \cong \mathbb{Z}^{2(g-1)}$ using the fact that $\widehat{\Sigma}_{g,\gamma}$ deformation retracts to S.

LEMMA 5.8. *We have a commutative diagram*

$$\begin{array}{ccccccccc} 1 & \longrightarrow & \mathbb{Z} & \longrightarrow & \pi_1(U\widehat{\Sigma}_{g,\gamma}) & \longrightarrow & \pi_1(\widehat{\Sigma}_{g,\gamma}) & \longrightarrow & 1 \\ & & \downarrow \cong & & \downarrow \psi & & \downarrow \overline{\psi} & & \\ 1 & \longrightarrow & \mathbb{Z} & \longrightarrow & K_g & \longrightarrow & \mathbb{Z}^{2(g-1)} & \longrightarrow & 1 \end{array} \tag{14}$$

where $\overline{\psi}$ is the abelianization map.

PROOF. We will use the notation $\phi_{w,q}$ introduced in §5.1. The surjection $\mathrm{Mod}_{g,\gamma} \to (\mathrm{Mod}_g)_\gamma$ from (10) takes $T_{\partial_1} \in \mathrm{Mod}_{g,\gamma}$ to $T_\gamma \in \mathrm{Mod}_g$. Since $[\gamma] = b_g$, Lemma 3.4 says that T_γ acts on $\mathrm{H}_1(\Sigma_g;\mathbb{Z})$ as

$$h \mapsto h + \hat{i}(b_g, h)b_g.$$

This is exactly the element $\phi_{0,-1}$ of K_g, which generates the kernel of the exact sequence

$$1 \longrightarrow \mathbb{Z} \longrightarrow K_g \longrightarrow \mathbb{Z}^{2(g-1)} \longrightarrow 1.$$

It follows immediately that we have a commutative diagram like (14); all that remains to check is that the induced map $\overline{\psi} : \pi_1(\widehat{\Sigma}_{g,\gamma}) \to \mathbb{Z}^{2(g-1)}$ is the abelianization map.

Since $\pi_1(\widehat{\Sigma}_{g,\gamma})$ is generated by simple closed curves, it is enough to verify that if $\delta \in \pi_1(\widehat{\Sigma}_{g,\gamma})$ is a simple closed curve, then $\overline{\psi}(\delta) = [\delta]$. To do this, it is enough to find some lift $\tilde{\delta} \in \pi_1(U\widehat{\Sigma}_{g,\gamma})$ of δ and check that $\psi(\tilde{\delta}) = \phi_{[\delta],0}$, i.e. that

$$\psi(\tilde{\delta})(a_g) = [\delta] + a_g.$$

This is immediate from Figure 5.2. $\square$

COROLLARY 5.9. *The homomorphism $\psi : \mathrm{Mod}_{g,\gamma} \to (\mathrm{Sp}_{2g}(\mathbb{Z}))_{b_g}$ is surjective.*

PROOF. Immediate from Lemmas 5.7 and 5.8. $\square$

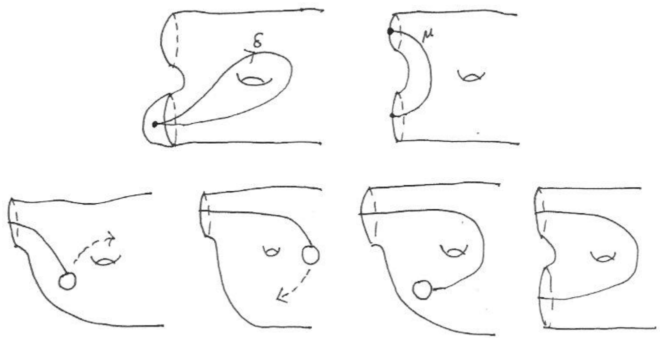

FIGURE 5.2. The bottom figures illustrate the effect of dragging the boundary component around the simple closed curve δ. The arc μ is chosen such that when the two boundary components are glued together, the two endpoints of μ match up to form an oriented simple closed curve homologous to a_g. As shown, dragging the boundary component around δ replaces μ with an arc homologous to $[\mu] + [\delta]$.

5.3. Stabilizers in the Torelli group

We finally discuss the Torelli group. Like in §5.2, fix some $g \geqslant 2$ and some oriented simple closed nonseparating curve γ on Σ_g.

Torelli on the cut-open surface. Let $\mathcal{I}_{g,\gamma}$ be the kernel of the symplectic representation $\psi : \mathrm{Mod}_{g,\gamma} \to (\mathrm{Sp}_{2g}(\mathbb{Z}))_{[\gamma]}$ discussed in §5.2. We then have the following. Let $\{\partial_1, \partial_2\}$ be the boundary components of $\Sigma_{g,\gamma}$.

LEMMA 5.10. *There is a short exact sequence*

$$1 \longrightarrow \mathbb{Z} \longrightarrow \mathcal{I}_{g,\gamma} \longrightarrow (\mathcal{I}_g)_\gamma \longrightarrow 1.$$

PROOF. An immediate consequence of Lemma 1.20. □

Semidirect product decomposition. We now come to the following theorem, which is the main result of this chapter. Versions of it were originally proved by van den Berg [**vdB03**, Proposition 2.4.1] and Putman [**Put07**, Theorem 4.1].

THEOREM 5.11. *Fix $g \geqslant 2$, let γ be an oriented simple closed nonseparating curve on Σ_g, and let S be a γ-splitting surface in $\Sigma_{g,\gamma}$. Letting $\widehat{\Sigma}_{g,\gamma}$ be the result of gluing a disc to one of the boundary components of $\Sigma_{g,\gamma}$ and $\pi = \pi_1(\widehat{\Sigma}_{g,\gamma})$, we then have a decomposition*

$$\mathcal{I}_{g,\gamma} = [\pi, \pi] \rtimes \mathcal{I}(S).$$

PROOF. Lemma 1.21 says that there is a decomposition

$$\mathrm{Mod}_{g,\gamma} = \pi_1(U\widehat{\Sigma}_{g,\gamma}) \rtimes \mathrm{Mod}(S).$$

Lemma 5.7 then implies that

$$\mathcal{I}_{g,\gamma} = \ker(\psi|_{\pi_1(U\widehat{\Sigma}_{g,\gamma})}) \rtimes \ker(\psi|_{\mathrm{Mod}(S)}).$$

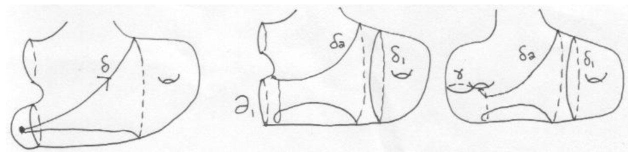

FIGURE 5.3. On the left is a curve $\delta \in [\pi, \pi]$ that can be realized by a simple closed separating curve. The next figure shows a particular lift $\tilde{\delta} = T_{\delta_1} T_{\delta_2}^{-1}$ of δ to $\pi_1(\widehat{\Sigma}_{g,\gamma})$; in fact, this is the element Push_δ discussed in §1.4. We claim that $T_{\delta_1} T_{\delta_2}^{-1} T_{\partial_1} \in \mathcal{I}_{g,\gamma}$. This is equivalent to saying that $T_{\delta_1} T_{\delta_2}^{-1} T_{\partial_1}$ maps to an element of $\mathcal{I}_g$ when ∂_1 and ∂_2 are glued back together. As is shown in the right most figure, the image of $T_{\delta_1} T_{\delta_2}^{-1} T_{\partial_1}$ in Mod_g is $T_{\delta_1} T_{\delta_2}^{-1} T_\gamma$, which lies in $\mathcal{I}_g$ since T_{δ_1} is a separating twist and $T_{\delta_2}^{-1} T_\gamma = T_\gamma T_{\delta_2}^{-1}$ is a bounding pair map. The T_{∂_1} is necessary here since T_{δ_2} is *not* a separating twist in $\mathcal{I}_g$ even though δ_2 does separate $\Sigma_{g,\gamma}$. The problem is that δ_2 is not nullhomologous in $\Sigma_{g,\gamma}$.

By definition we have $\ker(\psi|_{\text{Mod}(S)}) = \mathcal{I}(S)$. Also, Lemma 5.8 implies that the projection map $\pi_1(U\widehat{\Sigma}_{g,\gamma}) \to \pi$ takes $\ker(\psi|_{\pi_1(U\widehat{\Sigma}_{g,\gamma})})$ isomorphically onto the kernel of the abelianization map $\pi \to \mathbb{Z}^{2(g-1)}$, i.e. onto $[\pi, \pi]$. The theorem follows. □

Examples of elements. Let the notation be as in Theorem 5.11. The group $[\pi, \pi]$ is embedded in $\pi_1(\widehat{\Sigma}_{g,\gamma})$ in a somewhat complicated way. Consider some curve $\delta \in [\pi, \pi]$. The proof of Theorem 5.11 shows that the associated element of $\mathcal{I}_{g,\gamma}$ can be obtained as follows. Let

$$\tilde{\delta} \in \pi_1(U\widehat{\Sigma}_{g,\gamma}) \subset \text{Mod}_{g,\gamma}$$

be any lift of

$$\delta \in \pi = \pi_1(\widehat{\Sigma}_{g,\gamma}).$$

Then there exists a unique $k \in \mathbb{Z}$ such that $\tilde{\delta} T_{\partial_1}^k \in \mathcal{I}_{g,\gamma}$; this is the element of $\mathcal{I}_{g,\gamma}$ corresponding to δ. This takes a particularly simple form when δ can be realized by a simple closed separating curve; see Figure 5.3.

Separating twists and bounding pair maps. We will say that an element of $\mathcal{I}_{g,\gamma}$ is a *separating twist* (resp. a *bounding pair map*) if it maps to a separating twist (resp. a bounding pair map) in $\mathcal{I}_g \subset \text{Mod}_g$ when the boundary components ∂_1 and ∂_2 of $\Sigma_{g,\gamma}$ are glued back together. We will also say that $T_{\partial_1} T_{\partial_2}^{-1} \in \mathcal{I}_{g,\gamma}$ is a bounding pair map even though it maps to the identity in $\mathcal{I}_{g,\gamma}$. In the example discussed in Figure 5.3, the element of $\mathcal{I}_{g,\gamma}$ corresponding to the element $\delta \in [\pi, \pi]$ that can be realized by a separating simple closed curve is the product of a separating twist and a bounding pair map in $\mathcal{I}_{g,\gamma}$.

CHAPTER 6

The genus 2 Torelli group

In this chapter, we use the complex of cycles from Chapter 4 to prove Theorem 3.14, which says that $\mathcal{I}_2$ is an infinite-rank free group. This theorem is due to Mess [**Mes92**], but the proof we give is due to Bestvina–Bux–Margalit [**BBM10**]. We begin in §6.1 by discussing the topology of the complex of cycles on a genus 2 surface. Next, in §6.2 we prove a special case of the main theorem of Bass–Serre theory. We finally Theorem 3.14 in §6.3. This proof depends on a lemma whose proof is postponed until §6.4.

6.1. The complex of cycles in genus 2

This section is devoted to the structure of the complex of reduced cycles on a genus 2 surface.

Global topology. We begin with the following lemma.

LEMMA 6.1. *Let $x \in \mathrm{H}_1(\Sigma_2;\mathbb{Z})$ be a primitive element. Then the complex $\mathcal{C}_x(\Sigma_2)$ is a tree.*

PROOF. Theorem 4.14 says that $\mathcal{C}_x(\Sigma_2)$ is contractible, so it is enough to prove that $\mathcal{C}_x(\Sigma_2)$ is 1-dimensional. Let γ be an oriented multicurve on Σ_2 such that $\mathcal{X}_x(\gamma)$ is nondegenerate and let the components of Σ_2 cut along γ be $R_1,\ldots,R_\ell$. Corollary 4.2 says that $\mathcal{X}_x(\gamma)$ is $(\ell-1)$-dimensional, so it is enough to show that $\ell \leqslant 2$. Clearly none of the R_i are closed surfaces. Also, none of the components of γ are nullhomotopic, so none of the R_i are homeomorphic to a 1-holed sphere. Finally, no two components of γ are homotopic to each other (ignoring orientations), so none of the R_i are homeomorphic to 2-holed spheres. The upshot is that $\chi(R_i) \leqslant -1$ for $1 \leqslant i \leqslant \ell$. For $1 \leqslant i < j \leqslant \ell$ the intersection $R_i \cap R_j$ is a collection of circles. Since $\chi(S^1) = 0$, we conclude that

$$-2 = \chi(\Sigma_2) = \chi(R_1) + \cdots + \chi(R_\ell) \leqslant -\ell;$$

i.e. that $\ell \leqslant 2$, as desired. □

Realizing homology classes by multicurves. Our next goal is to study the action of the genus 2 Torelli group on the complex of reduced cycles. Our main result (Lemma 6.7 below) says that the quotient of the complex of reduced cycles on a genus 2 surface by the Torelli group is also a tree. We begin by describing which collections of homology classes can be realized by nonseparating multicurves. Here we work on an arbitrary genus g surface. A collection $\{v_1,\ldots,v_k\}$ of elements of $\mathrm{H}_1(\Sigma_g;\mathbb{Z})$ is *unimodular* if it is a basis for a direct summand of $\mathrm{H}_1(\Sigma_g;\mathbb{Z}) \cong \mathbb{Z}^{2g}$ and *isotropic* if $\hat{i}(v_i,v_j) = 0$ for all $1 \leqslant i,j \leqslant k$. See Figure 6.1 for a picture of the curves in the lemma below.

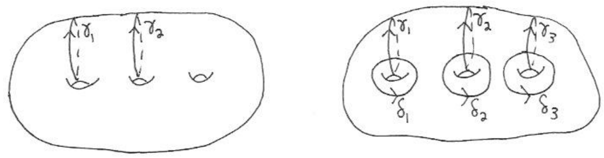

FIGURE 6.1. On the left is a collection $\{\gamma_1, \gamma_2\}$ of disjoint oriented simple closed curves whose union does not separate the surface. On the right is how to complete this to a geometric symplectic basis.

LEMMA 6.2. *For some $g \geqslant 1$, let $\{v_1, \ldots, v_k\}$ be a collection of distinct elements of $\mathrm{H}_1(\Sigma_g; \mathbb{Z})$. Then there exist disjoint oriented simple closed curves $\{\gamma_1, \ldots, \gamma_k\}$ on Σ_g such that $\gamma_1 \cup \cdots \cup \gamma_k$ does not separate Σ_g and such that $[\gamma_i] = v_i$ for $1 \leqslant i \leqslant k$ if and only if the v_i are unimodular and isotropic.*

PROOF. Assume first that such curves $\{\gamma_1, \ldots, \gamma_k\}$ exist. Since the γ_i are disjoint and $\gamma_1 \cup \cdots \cup \gamma_k$ does not separate Σ_g, we can find oriented simple closed curves $\gamma_{k+1}, \ldots, \gamma_g, \delta_1, \ldots, \delta_g$ on Σ_g such that $\{\gamma_1, \delta_1, \ldots, \gamma_g, \delta_g\}$ is a geometric symplectic basis (see Figure 6.1). This implies that $\{[\gamma_1], [\delta_1], \ldots, [\gamma_g], [\delta_g]\}$ is a symplectic basis for $\mathrm{H}_1(\Sigma_g; \mathbb{Z})$. It follows immediately that $\{[\gamma_1], \ldots, [\gamma_k]\} = \{v_1, \ldots, v_k\}$ is a set of elements which is unimodular and isotropic.

Now assume that $\{v_1, \ldots, v_k\}$ is a set of elements of $\mathrm{H}_1(\Sigma_g; \mathbb{Z})$ which is unimodular and isotropic. Since any symplectic basis can be realized by a geometric symplectic basis (Proposition 2.10), it is enough to prove that there exist $v_{k+1}, \ldots, v_g, w_1, \ldots, w_g \in \mathrm{H}_1(\Sigma_g; \mathbb{Z})$ such that $\{v_1, w_1, \ldots, v_g, w_g\}$ is a symplectic basis. Let $X \subset \mathrm{H}_1(\Sigma_g; \mathbb{Z})$ be the span of the v_i. Since the v_i are unimodular, we can find $X' \subset \mathrm{H}_1(\Sigma_g; \mathbb{Z})$ such that $\mathrm{H}_1(\Sigma_g; \mathbb{Z}) = X \oplus X'$. For $1 \leqslant i \leqslant k$, define a linear map $\phi_i : \mathrm{H}_1(\Sigma_g; \mathbb{Z}) \to \mathbb{Z}$ via the formulas

$$\phi_i|_{X'} = 0 \quad \text{and} \quad \phi_i(v_i) = 1 \quad \text{and} \quad \phi_i(v_j) = 0 \quad \text{for } j \neq i.$$

Since $\hat{i}(\cdot, \cdot)$ is a symplectic form (Lemma 2.1), we can find elements $w_1, \ldots, w_k \in \mathrm{H}_1(\Sigma_g; \mathbb{Z})$ such that for $1 \leqslant i \leqslant k$, we have $\phi_i(u) = \hat{i}(u, w_i)$ for all $u \in \mathrm{H}_1(\Sigma_g; \mathbb{Z})$. Let Y be the span of $\{v_1, w_1, \ldots, v_k, w_k\}$. Define a homomorphism $\psi : \mathrm{H}_1(\Sigma_g; \mathbb{Z}) \to \mathbb{Z}^{2k}$ via the formula

$$\psi(u) = (\hat{i}(v_1, u), \hat{i}(w_1, u), \ldots, \hat{i}(v_k, u), \hat{i}(w_k, u)).$$

Clearly ψ takes Y isomorphically onto $\mathbb{Z}^{2k}$. Defining $Z = \ker(\psi)$, we obtain that $\mathrm{H}_1(\Sigma_g; \mathbb{Z}) \cong Y \oplus Z$. It is easy to see that $\hat{i}(\cdot, \cdot)$ restricts to a symplectic form on Z, so we can find a symplectic basis $\{v_{k+1}, w_{k+1}, \ldots, v_g, w_g\}$ for Z. The set $\{v_1, w_1, \ldots, v_g, w_g\}$ is then the desired symplectic basis for $\mathrm{H}_1(\Sigma_g; \mathbb{Z})$. □

Vertices of quotient. We continue to work on an arbitrary genus g surface. If $x \in \mathrm{H}_1(\Sigma_g; \mathbb{Z})$ is a primitive element, then Lemma 4.3 says that the vertices of $\mathcal{C}_x(\Sigma_g)$ are exactly the points $c_1\gamma_1 + \cdots + c_k\gamma_k$, where $\gamma = \gamma_1 \cup \cdots \cup \gamma_k$ is an oriented multicurve that does not separate Σ_g and $c_1, \ldots, c_k$ are positive integers such that

$$x = c_1[\gamma_1] + \cdots + c_k[\gamma_k].$$

Proposition 3.24 implies that two vertices $c_1\gamma_1 + \cdots + c_k\gamma_k$ and $c_1'\gamma_1' + \cdots + c_{k'}'\gamma_{k'}'$ of $\mathcal{C}_x(\Sigma_g)$ are in the same $\mathcal{I}_g$-orbit if and only if $k = k'$ and $[\gamma_i] = [\gamma_i']$ and $c_i = c_i'$ for

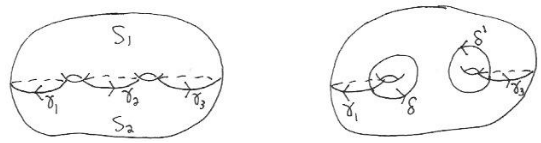

FIGURE 6.2. On the left are three nonseparating simple closed curves $\{\gamma_1, \gamma_2, \gamma_3\}$ on Σ_2 that are pairwise disjoint and nonhomotopic. Their union separates Σ_2 into two subsurfaces S_1 and S_2. On the right are the two curves δ and δ' constructed in the proof of Lemma 6.5.

$1 \leqslant i \leqslant k$. Setting $v_i = [\gamma_i]$ for $1 \leqslant i \leqslant k$, we denote this orbit by the formal symbol $c_1[\![v_1]\!] + \cdots + c_k[\![v_k]\!]$. We will say that k is the *size* of the vertex $c_1[\![v_1]\!] + \cdots + c_k[\![v_k]\!]$. Combining these observations with Lemma 6.2, we deduce the following.

LEMMA 6.3. *For $g \geqslant 1$, fix some primitive element $x \in \mathrm{H}_1(\Sigma_g; \mathbb{Z})$. Then the vertices of $\mathcal{C}_x(\Sigma_g)/\mathcal{I}_2$ are the formal symbols $c_1[\![v_1]\!] + \cdots + c_k[\![v_k]\!]$, where $c_1, \ldots, c_k$ are positive integers, the set $\{v_1, \ldots, v_k\}$ is a unimodular and isotropic set of elements of $\mathrm{H}_1(\Sigma_g; \mathbb{Z})$, and $x = c_1 v_1 + \cdots + c_k v_k$.*

Edges in genus 2. We now describe the edges of the quotient of the complex of reduced cycles by the Torelli group in genus 2.

LEMMA 6.4. *Let $x \in \mathrm{H}_1(\Sigma_2; \mathbb{Z})$ be a primitive element and let $c_1[\![v_1]\!] + c_2[\![v_2]\!]$ be a vertex of $\mathcal{C}_x(\Sigma_2)/\mathcal{I}_2$ of size 2. Order the v_i such that $c_1 \geqslant c_2$. Then there are exactly three edges containing $c_1[\![v_1]\!] + c_2[\![v_2]\!]$. Their other endpoints are $(c_1 - c_2)[\![v_1]\!] + c_2[\![v_1 + v_2]\!]$ and $(c_1 + c_2)[\![v_2]\!] + c_1[\![v_1 - v_2]\!]$ and $(c_1 + c_2)[\![v_1]\!] + c_2[\![v_2 - v_1]\!]$.*

Before we prove Lemma 6.4, we need two auxiliary lemmas.

LEMMA 6.5. *Let $\gamma_1 \cup \gamma_2 \cup \gamma_3$ be an oriented multicurve on Σ_2. Assume that none of the γ_i separate Σ_2. Then the symplectic representation $\psi : \mathrm{Mod}_2 \to \mathrm{Sp}_4(\mathbb{Z})$ restricts to an isomorphism from the stabilizer subgroup*

$$\Gamma = \{f \in \mathrm{Mod}_2 \mid f(\gamma_i) = \gamma_i \text{ for } 1 \leqslant i \leqslant 3\}$$

to the stabilizer subgroup

$$G = \{f \in \mathrm{Sp}_4(\mathbb{Z}) \mid f([\gamma_i]) = \gamma_i \text{ for } 1 \leqslant i \leqslant 3\}.$$

PROOF. An Euler characteristic argument similar to the one that appeared in the proof of Lemma 6.1 shows that $\gamma_1 \cup \gamma_2 \cup \gamma_3$ separates Σ_2 into two subsurfaces S_1 and S_2, each of which is homeomorphic to a 3-holed sphere (see Figure 6.2). For $j = 1, 2$ we have $\mathrm{Mod}(S_j) \cong \mathbb{Z}^3$ with generators $\{T_{\gamma_1}, T_{\gamma_2}, T_{\gamma_3}\}$ (this is an easy exercise; see [**FM12**, §3.6.4] for details). It follows that $\Gamma \cong \mathbb{Z}^3$ with generators $\{T_{\gamma_1}, T_{\gamma_2}, T_{\gamma_3}\}$.

As in Figure 6.2, pick disjoint oriented simple closed curves δ and δ' on Σ_2 such that δ (resp. δ') intersects γ_1 and γ_2 (resp. γ_2 and γ_3) each once. Orienting δ and δ' appropriately, we can arrange for $\{[\gamma_1], [\delta], [\gamma_3], [\delta']\}$ to be a symplectic basis for $\mathrm{H}_1(\Sigma_2; \mathbb{Z})$. To keep our notation from getting out of hand, define

$$a_1 = [\gamma_1] \quad \text{and} \quad b_1 = [\delta] \quad \text{and} \quad a_2 = [\gamma_3] \quad \text{and} \quad b_2 = [\delta'].$$

Consider $f \in G$. By definition, we have $f(a_1) = a_1$ and $f(a_2) = a_2$. Since $\hat{i}(f(a_1), f(b_1)) = 1$ and $\hat{i}(f(a_2), f(b_1)) = 0$, we must have

$$f(b_1) = b_1 + c_1 a_1 + c_2 a_2$$

for some $c_1, c_2 \in \mathbb{Z}$. Similarly, we must have

$$f(b_2) = b_2 + d_1 a_1 + d_2 a_2$$

for some $d_1, d_2 \in \mathbb{Z}$. We then have

$$0 = \hat{i}(f(b_1), f(b_2)) = -d_1 + c_2,$$

i.e. $d_1 = c_2$. Define $\phi : G \to \mathbb{Z}^3$ by $\phi(f) = (c_1, c_2, d_2)$. It is easy to see that ψ is a homomorphism, and by construction the kernel of ϕ is trivial, i.e. ϕ is an isomorphism.

It is enough now to prove that

$$\{\phi \circ \psi(T_{\gamma_1}), \phi \circ \psi(T_{\gamma_2}), \phi \circ \psi(T_{\gamma_3})\}$$

is a basis for $\mathbb{Z}^3$. First, we have

$$T_{\gamma_1}(b_1) = b_1 + a_1 \quad \text{and} \quad T_{\gamma_1}(b_2) = b_2,$$

so $\phi \circ \psi(T_{\gamma_1}) = (1, 0, 0)$. Similarly, we have $\phi \circ \psi(T_{\gamma_3}) = (0, 0, 1)$. Finally, as is clear from Figure 6.2 we have $[\gamma_2] = e_1 a_1 + e_2 a_2$ for some $e_1, e_2 \in \{-1, 1\}$. It follows that

$$T_{\gamma_2}(b_1) = b_1 + a_1 + e_1 e_2 a_2 \quad \text{and} \quad T_{\gamma_2}(b_2) = b_2 + e_1 e_2 a_1 + a_2,$$

so $\phi \circ \psi(T_{\gamma_2}) = (1, e_1 e_2, 1)$. The lemma follows. □

LEMMA 6.6. *Let $\gamma_1 \cup \gamma_2 \cup \gamma_3$ and $\gamma_1' \cup \gamma_2' \cup \gamma_3'$ be two oriented multicurves on Σ_2. Assume that none of the γ_i or γ_i' separate Σ_2 and that $[\gamma_i] = [\gamma_i']$ for $1 \leqslant i \leqslant 3$. Then there exists a unique $f \in \mathcal{I}_2$ such that $f(\gamma_i) = \gamma_i'$ for $1 \leqslant i \leqslant 3$.*

PROOF. Just like in the proof of Lemma 6.5, the multicurve $\gamma_1 \cup \gamma_2 \cup \gamma_3$ separates Σ_2 into two subsurfaces S_1 and S_2, each of which is homeomorphic to a 3-holed sphere (see Figure 6.2). Order the S_i such that S_1 lies to the right of γ_1. Similarly, $\gamma_1' \cup \gamma_2' \cup \gamma_3'$ separates Σ_2 into two subsurfaces S_1' and S_2', each of which is homeomorphic to a 3-holed sphere, and we order the S_i' such that S_1' lies to the right of γ_1'. The subsurface S_1 and S_1' give homologies showing that

$$[\gamma_1] + e_2[\gamma_2] + e_3[\gamma_3] = 0 \quad \text{and} \quad [\gamma_1'] + e_2'[\gamma_2'] + e_3'[\gamma_3'] = 0$$

for some $e_2, e_3, e_2', e_3' \in \{-1, 1\}$. Since $[\gamma_2] = [\gamma_2']$ and $[\gamma_3] = [\gamma_3']$, we see that $e_2 = e_2'$ and $e_3 = e_3'$. In other words, for $1 \leqslant i \leqslant 3$ the orientations of γ_i and the boundary of S_2 agree/disagree exactly when the the orientations of γ_i' and the boundary of S_2' agree/disagree. The classification of surfaces trick thus say that there exists some mapping class $f' \in \mathrm{Mod}_2$ such that $f'(\gamma_i) = \gamma_i'$ for $1 \leqslant i \leqslant 3$.

Any $f \in \mathrm{Mod}_2$ satisfying $f(\gamma_i) = \gamma_i'$ for $1 \leqslant i \leqslant 3$ can be written $f = f'h$, where $h \in \mathrm{Mod}_2$ is a mapping class that fixes γ_i for $1 \leqslant i \leqslant 3$. By construction, the image $\overline{f}'$ of f' in $\mathrm{Sp}_4(\mathbb{Z})$ fixes $[\gamma_i]$ for $1 \leqslant i \leqslant 3$. Lemma 6.5 therefore says that there exists a unique $h \in \mathrm{Mod}_2$ that fixes γ_i for $1 \leqslant i \leqslant 3$ such that the image of h in $\mathrm{Sp}_4(\mathbb{Z})$ is $(\overline{f}')^{-1}$. This implies that there exists a unique $f \in \mathrm{Mod}_2$ satisfying $f(\gamma_i) = \gamma_i'$ for $1 \leqslant i \leqslant 3$ such that the image of f in $\mathrm{Sp}_4(\mathbb{Z})$ is trivial, i.e. such that $f \in \mathcal{I}_2$. □

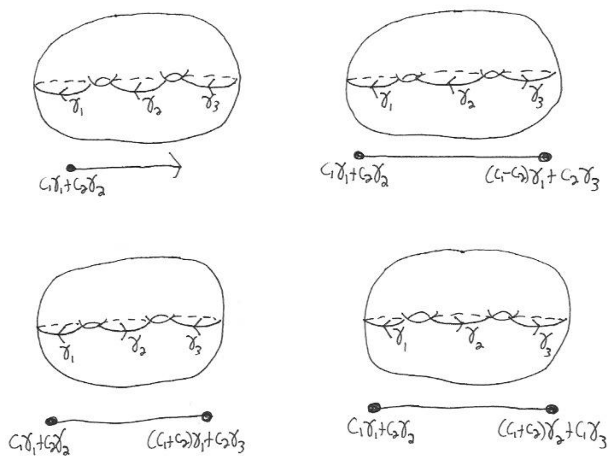

FIGURE 6.3. Realizing the four possible homology classes for the third curve γ_3 in the proof of Lemma 6.4. As in that proof, we have $x = c_1[\gamma_1] + c_2[\gamma_2]$. Below each configuration of curves is the associated cell of $\widehat{\mathcal{C}}_x(\Sigma_2)$. The upper left hand configuration gives a non-reduced cell.

PROOF OF LEMMA 6.4. Let $\gamma_1 \cup \gamma_2$ be an oriented multicurve such that $v_1 = [\gamma_1]$ and $v_2 = [\gamma_2]$. Using Corollary 4.2, the edges of $\mathcal{C}_x(\Sigma_2)/\mathcal{I}_2$ that contain $c_1[\![v_1]\!] + c_2[\![v_2]\!]$ are exactly the $\mathcal{I}_2$-orbits of nondegenerate cells $\mathcal{X}_x(\gamma)$ such that γ is a multicurve that separates Σ_2 into two components and contains $\gamma_1 \cup \gamma_2$. Such multicurves have three components $\gamma_1 \cup \gamma_2 \cup \gamma_3$, and Lemma 6.6 says that their $\mathcal{I}_2$-orbits are determined by $[\gamma_3] \in \mathrm{H}_1(\Sigma_2;\mathbb{Z})$. Moreover, since $\gamma_1 \cup \gamma_2 \cup \gamma_3$ separates Σ_2 into two components, we must have $[\gamma_3] = e_1[\gamma_1] + e_2[\gamma_2]$ for some signs $e_1, e_2 \in \{-1, 1\}$ (just like in Figure 6.2). As is shown in Figure 6.3, all four possibilities are actually realized. The cell with $[\gamma_3] = -[\gamma_1] - [\gamma_2]$ is not reduced, and the other three edges are exactly the edges described in the statement of the lemma. □

The quotient is a tree. We finally come to the following lemma.

LEMMA 6.7. *Let $x \in \mathrm{H}_1(\Sigma_2;\mathbb{Z})$ be a primitive element. Then $\mathcal{C}_x(\Sigma_2)/\mathcal{I}_2$ is a tree.*

PROOF. By Lemma 6.3, the vertices of $\mathcal{C}_x(\Sigma_2)/\mathcal{I}_2$ are exactly $[\![x]\!]$ and $c_1[\![v_1]\!] + c_2[\![v_2]\!]$, where $v_1, v_2 \in \mathrm{H}_1(\Sigma_2;\mathbb{Z})$ are unimodular and isotropic and $c_1, c_2 \in \mathbb{Z}$ are positive and satisfy

$$x = c_1 v_1 + c_2 v_2.$$

Say that the *height* of $[\![x]\!]$ is 1 and that the height of $c_1[\![v_1]\!] + c_2[\![v_2]\!]$ is $c_1 + c_2$. There is thus a unique vertex $[\![x]\!]$ of height 1. Also, examining the edges given in Lemma 6.4 we see that for a size 2 vertex $c_1[\![v_1]\!] + c_2[\![v_2]\!]$ of $\mathcal{C}_x(\Sigma_2)/\mathcal{I}_2$, there is a unique edge coming out of $c_1[\![v_1]\!] + c_2[\![v_2]\!]$ which ends at a vertex of smaller height. The other two edges coming out of $c_1[\![v_1]\!] + c_2[\![v_2]\!]$ end at vertices of larger height. This immediately implies that $\mathcal{C}_x(\Sigma_2)/\mathcal{I}_2$ is a tree. Indeed, assume that $p_1, \ldots, p_k$

are vertices of $\mathcal{C}_x(\Sigma_2)/\mathcal{I}_2$ such that

$$p_0 - p_2 - \cdots - p_k$$

is an embedded loop. Thus $p_0 = p_k$ and $p_i \neq p_j$ for $0 \leqslant i, j \leqslant k$ with $\{i, j\} \neq \{0, k\}$. Choose $1 \leqslant \ell \leqslant k$ such that the height of p_ℓ is maximal among the heights of the p_i. Then the heights of $p_{\ell+1}$ and $p_{\ell-1}$ (indices taken modulo k) are strictly less than the height of p_ℓ, which is impossible. □

6.2. A little Bass–Serre theory

Bass–Serre theory is the study of group actions on trees. We will need a very special case of it. Two classic sources for the general case are the book [**Ser80**] by Serre and the long survey [**SW79**] by Scott–Wall. Our approach is close to the combinatorial techniques of [**Ser80**]; the paper [**SW79**] is more topological. All group actions on simplicial complexes in this section are assumed to be simplicial.

Strict fundamental domains. Consider a group G acting on a simplicial complex X. A *strict fundamental domain* for the action of G on X is a subcomplex D of X such that for all simplices σ of X, there is a unique simplex σ' of D in the G-orbit of σ. If D is a strict fundamental domain for the action of G on X, then the following two things hold.

- The group G acts *without rotation* on X, that is, for all simplices σ of X, the stabilizer subgroup G_σ of σ stabilizes σ pointwise. This implies in particular that X/G is a cell-complex whose cells are exactly the G-orbits of simplices in X. We remark that the quotient spaces of groups acting without rotations on simplicial complexes need not be simplicial complexes in general; see the remark after the second example below.
- The projection map $X \to X/G$ restricts to an isomorphism $D \cong X/G$ of cell complexes.

Here is an example and a non-example. In both of these examples, $\mathbb{R}$ is triangulated by placing a vertex at each integer.

Example 6.8. The infinite dihedral group

$$D_\infty = \langle s, t \mid s^2 = 1,\ sts^{-1} = t^{-1} \rangle$$

acts on $\mathbb{R}$ via the formulas

$$s(x) = -x \qquad \text{and} \qquad t(x) = x + 2$$

for $x \in \mathbb{R}$. The edge $[0, 1]$ is a strict fundamental domain for this action. Observe that the stabilizers of the vertices 0 and 1 are both isomorphic to $\mathbb{Z}/2$, with the former generated by s and the latter generated by ts. Also, the stabilizer of the entire edge $[0, 1]$ is trivial.

Example 6.9. The group $\mathbb{Z} = \langle s \mid \rangle$ acts on $\mathbb{R}$ via the formula

$$s(x) = x + 1$$

for $x \in \mathbb{R}$. There is no strict fundamental domain for this action. Indeed, the quotient $\mathbb{R}/\mathbb{Z}$ is a circle S^1, and no subcomplex of $\mathbb{R}$ is homeomorphic to S^1.

Remark 6.10. The above action of $\mathbb{Z}$ on $\mathbb{R}$ is without rotations. While the quotient $\mathbb{R}/\mathbb{Z}$ is a cell complex with a single vertex and a single edge, it is *not* a simplical complex.

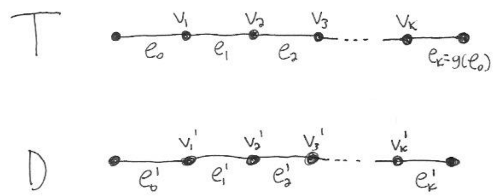

FIGURE 6.4. Illustration of Step 1 in the proof of Theorem 6.12, which asserts that ρ is surjective. On the top is the edge-path in the tree T, and on the bottom is the projection of this path to the strict fundamental domain D.

Existence of strict fundamental domains. The following lemma characterizes which group actions on trees have strict fundamental domains.

LEMMA 6.11. *Let G be a group acting without rotations on a tree T. There exists a strict fundamental domain for the action of G on T if and only if T/G is a tree.*

PROOF. If a strict fundamental domain D exists, then as noted above we have $D \cong T/G$. Since $D \cong T/G$ is connected and connected subgraphs of trees are themselves trees, we deduce that T/G is a tree.

Conversely, assume that T/G is a tree. Let $\mathcal{D}$ be the set of all connected subgraphs of T that map isomorphically into T/G, partially ordered by inclusion. Clearly $\mathcal{D}$ satisfies the conditions of Zorn's lemma, so it contains a maximal element D. If D is not a strict fundamental domain, then there must exist an edge $\overline{e}$ of T/G such that one vertex of $\overline{e}$ lies in the image of D and the other vertex does not (this is where we use the fact that T/G is a tree). We can lift $\overline{e}$ to an edge e of T such that one vertex of e lies in D and the other does not. Then $D \cup e \in \mathcal{D}$, contradicting the maximality of D. □

Decompositions from strict fundamental domains. The following theorem is a special case of the "fundamental theorem of Bass–Serre theory". Applied to Example 6.9 above, it says that $D_\infty = \mathbb{Z}/2 * \mathbb{Z}/2$.

THEOREM 6.12. *Let G be a group acting on a tree T. Assume that $D \subset T$ is a strict fundamental domain for the action of G and that for all edges e of T, the stabilizer subgroup G_e is trivial. Then*

$$G = \mathop{\Huge *}_{v \in D^{(0)}} G_v.$$

REMARK 6.13. If $G_e \neq 1$ for some edge e with endpoints v and v', then one has to identify the images of G_e in G_v and $G_{v'}$. See [**Ser80**] and [**SW79**] for more details.

PROOF OF THEOREM 6.12. Define

$$\Gamma = \mathop{\Huge *}_{v \in D^{(0)}} G_v.$$

There is a natural projection map $\rho : \Gamma \to G$.

STEP 1. *The map ρ is surjective.*

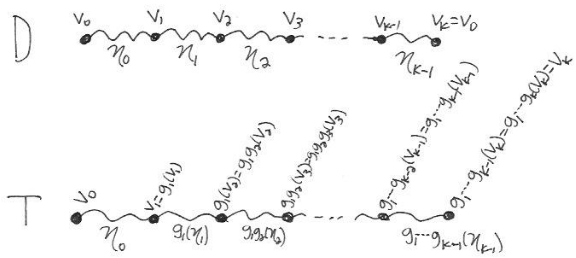

FIGURE 6.5. Illustration of Step 2 in the proof of Theorem 6.12, which asserts that ρ is injective. On the top is the loop that we construct in the strict fundamental domain D, and on the bottom is the resulting loop in T. The loop in D on the top is not locally injective, but we prove that the loop in T on the bottom is.

This step in the proof is illustrated in Figure 6.4. Consider $g \in G \backslash \{1\}$. Fixing an edge e_0 of D, there exists a path in T starting with the edge e_0 and ending with the edge $g(e_0)$. Since $G_{e_0} = 1$, this path has at least 2 edges in it. Let its edges be

$$e_0 - e_1 - \cdots - e_k = g(e_0). \tag{15}$$

For each $1 \leqslant i \leqslant k$, there exists a unique edge e_i' of D and some $h_i \in G$ such that $e_i = h_i(e_i')$. Choose the h_i such that $h_0 = 1$ and $h_k = g$. The edges

$$e_0' - e_1' - \cdots - e_k' \tag{16}$$

form a path in D; indeed, identifying D with T/G, this is the projection of (15) to T/G. For $0 < i \leqslant k$, let v_i' be the vertex of D that is shared by e_i' and e_{i-1}' in (16). Letting v_i be the vertex of T that is shared by e_i and e_{i-1} in (15), we have $h_i(v_i') = h_{i-1}(v_i') = v_i$. It follows that $h_{i-1}^{-1}h_i(v_i') = v_i'$, so $h_{i-1}^{-1}h_i \in \psi(\Gamma)$. Using the fact that $h_0 = 1$, we have that

$$g = h_k = (h_0^{-1}h_1)(h_1^{-1}h_2)\cdots(h_{k-1}^{-1}h_k) \in \psi(\Gamma),$$

as desired.

STEP 2. *The map ρ is injective.*

This step in the proof is illustrated in Figure 6.5. Assume that ρ is not injective, so there exists some nontrivial $w \in \ker(\rho)$. As notation, for $v \in D^{(0)}$ and $g \in G_v$, we will denote by g_v the associated element of Γ. Write

$$w = (g_1)_{v_1} \cdots (g_k)_{v_k},$$

where the g_i and v_i satisfy the following conditions.

- For $1 \leqslant i \leqslant k$, we have $v_i \in D^{(0)}$ and $g_i \in G_{v_i} \backslash \{1\}$.
- For $1 \leqslant i < k$, we have $v_i \neq v_{i+1}$.

Choose w such that k is as small as possible. Clearly $k > 1$. Also, if $v_1 = v_k$, then w is conjugate to

$$(g_k^{-1}g_1)_{v_1}(g_2)_{v_2}\cdots(g_{k-1})_{v_{k-1}}.$$

Since $\ker(\rho)$ is normal, this contradicts the minimality of k. We deduce that $v_1 \neq v_k$. Set $v_0 = v_k$.

For $0 \leqslant i < k$, let η_i be a simple (i.e. embedded) edge-path path in D from v_i to v_{i+1}. For $1 \leqslant i \leqslant k$, the fact that $g_i \in G_{v_i}$ implies that the terminal point of

η_{i-1} and the starting point of $g_i(\eta_i)$ are both v_i. This implies that the terminal point of $g_1 \cdots g_{i-1}(\eta_{i-1})$ and the initial point of $g_1 \cdots g_{i-1} g_i(\eta_i)$ are the same. We thus have a path

$$\eta_0 - g_1(\eta_1) - g_1 g_2(\eta_2) - \cdots - g_1 g_2 \cdots g_{k-1}(\eta_{k-1}) \tag{17}$$

in T. The final point of $g_1 g_2 \cdots g_{k-1}(\eta_{k-1})$ is

$$g_1 g_2 \cdots g_{k-1}(v_k) = g_1 g_2 \cdots g_{k-1} g_k(v_k) = v_k = v_0;$$

here we are using the fact that $g_k(v_k) = v_k$ and $g_1 \cdots g_k = 1$. The path (17) is thus a closed path. The space T is a tree, so the closed path (17) must not be locally injective. Since each η_i is simple and nontrivial, this implies that for some $0 \leqslant i < k-1$ the final edge of $g_1 \cdots g_i(\eta_i)$ must be the same as the initial edge of $g_1 \cdots g_{i+1}(\eta_{i+1})$. Equivalently, the final edge of η_i must be the same as the initial edge of $g_{i+1}(\eta_{i+1})$.

Let e be the final edge of η_i and let e' be the initial edge of η_{i+1}. We thus have $e' \in D^{(1)}$ and $g_{i+1}(e') = e \in D^{(1)}$. Since D is a strict fundamental domain, we must have $e = e'$ and $g_{i+1}(e) = e$. But $G_e = 1$, so $g_{i+1} = 1$, a contradiction. □

6.3. Mess's theorem

This section is devoted to the proof of Theorem 3.14, which asserts that $\mathcal{I}_2$ is an infinite-rank free group.

Separating splittings. In fact, we will prove a more precise result. Recall from §3.4 that if δ is a simple closed separating curve on Σ_g, then the *separating splitting* induced by δ is defined as follows. Let S_1 and S_2 be the subsurfaces of Σ_g obtained by cutting Σ_g along δ. The separating splitting induced by δ is then the unordered pair (U_1, U_2), where U_i is the image of $\mathrm{H}_1(S_i; \mathbb{Z})$ in $\mathrm{H}_1(\Sigma_g; \mathbb{Z})$. Each U_i is a symplectic subspace of $\mathrm{H}_1(\Sigma_g; \mathbb{Z})$, and the U_i are orthogonal with respect to the algebraic intersection form in the sense that $\hat{i}(u_1, u_2) = 0$ for $u_1 \in U_1$ and $u_2 \in U_2$. Corollary 3.28 says that two separating twists T_δ and $T_{\delta'}$ in $\mathcal{I}_g$ are conjugate in $\mathcal{I}_g$ if and only if δ and δ' induce the same separating splitting.

Improved result. The above implies that if $B \subset \mathcal{I}_2$ is a free basis, then B can contain at most one separating twist T_δ inducing a given separating splitting. Define $\mathcal{S}$ to be the set of all possible separating splittings of $\mathrm{H}_1(\Sigma_2; \mathbb{Z})$. More precisely, $\mathcal{S}$ consists of all unordered pairs (U_1, U_2), where the U_i are orthogonal 2-dimensional symplectic subspaces of $\mathrm{H}_1(\Sigma_2; \mathbb{Z})$ such that $\mathrm{H}_1(\Sigma_2; \mathbb{Z}) = U_1 \oplus U_2$. We then have the following strengthening of Theorem 3.14. It was originally proved by Mess in his thesis [**Mes92**].

THEOREM 6.14. *There exists a set C of simple closed separating curves on Σ_2 with the following two properties.*

- *$\mathcal{I}_2$ is a free group on the free basis $\{T_\delta \mid \delta \in C\}$.*
- *There exists a bijection $\phi : C \to \mathcal{S}$ such that for $\delta \in C$, the curve δ induces the separating splitting $\phi(\delta)$.*

We will prove Theorem 6.14 at the end of this section. The proof we will give is due to Bestvina–Bux–Margalit [**BBM10**].

Dividing up the splittings. Fix a primitive element $x \in \mathrm{H}_1(\Sigma_2; \mathbb{Z})$. In preparation for the proof of Theorem 6.14, we partition $\mathcal{S}$ into subsets labeled by the vertices of $\mathcal{C}_x(\Sigma_2)/\mathcal{I}_2$. Lemma 6.3 says that these vertices fall into the following two classes.

- The vertex $[\![x]\!]$. Define
$$\mathcal{S}_{[\![x]\!]} = \{(U_1, U_2) \in \mathcal{S} \mid x \in U_1\}.$$
- Size 2 vertices $c_1[\![v_1]\!] + c_2[\![v_2]\!]$. Define
$$\mathcal{S}_{c_1[\![v_1]\!]+c_2[\![v_2]\!]} = \{(U_1, U_2) \in \mathcal{S} \mid v_1 \in U_1,\ v_2 \in U_2\}.$$

We then have the following.

LEMMA 6.15. *Fix a primitive element $x \in \mathrm{H}_1(\Sigma_2; \mathbb{Z})$. Then $\mathcal{S}$ is the disjoint union of the sets $\mathcal{S}_v$ as v ranges over the vertices of $\mathcal{C}_x(\Sigma_2)/\mathcal{I}_2$.*

PROOF. Consider $(U_1, U_2) \in \mathcal{S}$. We must show that there exists a unique vertex v of $\mathcal{C}_x(\Sigma_2)/\mathcal{I}_2$ such that $(U_1, U_2) \in \mathcal{S}_v$. Since $\mathrm{H}_1(\Sigma_2; \mathbb{Z}) = U_1 \oplus U_2$, there is a unique expression $x = c_1 v_1 + c_2 v_2$ with the v_i and c_i as follows for $i = 1, 2$.

- $v_i \in U_i$ is a primitive element or 0.
- $c_i \in \mathbb{Z}$ satisfies $c_i \geqslant 0$. Also, $c_i = 0$ if and only if $v_i = 0$.

We then have $(U_1, U_2) \in \mathcal{S}_{c_1[\![v_1]\!]+c_2[\![v_2]\!]}$, and $c_1[\![v_1]\!] + c_2[\![v_2]\!]$ is the unique vertex with this property. We remark that $c_1[\![v_1]\!] + c_2[\![v_2]\!] = [\![x]\!]$ exactly when one of the c_i vanishes. ∎

Vertex stabilizers Continue to let $x \in \mathrm{H}_1(\Sigma_2; \mathbb{Z})$ be a fixed primitive element. Our next goal is to understand the stabilizers in $\mathcal{I}_2$ of the vertices of $\mathcal{C}_x(\Sigma_2)$. These are given by the following lemma.

LEMMA 6.16. *Let $x \in \mathrm{H}_1(\Sigma_2; \mathbb{Z})$ be a primitive element and let c be a vertex of $\mathcal{C}_x(\Sigma_2)$. Then there exists a set C_c of simple closed separating curve on Σ_2 with the following three properties.*

- *Each $\delta \in C_c$ is disjoint from the oriented multicurve on which c is supported.*
- *The stabilizer subgroup $(\mathcal{I}_2)_c$ is a free group on the free basis $\{T_\delta \mid \delta \in C_c\}$.*
- *Let $\bar{c}$ be the image of c in $\mathcal{C}_x(\Sigma_2)/\mathcal{I}_2$. Then there exists a bijection $\phi : C_c \to \mathcal{S}_{\bar{c}}$ such that for $\delta \in C_c$, the curve δ induces the separating splitting $\phi(\delta)$.*

Lemma 6.16 is proved below in §6.4. Its proof makes use of the results from Chapter 5 on stabilizers in Torelli of nonseparating simple closed curves.

The proof. We now prove Theorem 6.14.

PROOF OF THEOREM 6.14. Fix some primitive element $x \in \mathrm{H}_1(\Sigma_2; \mathbb{Z})$. The group $\mathcal{I}_2$ acts without rotations on the complex $\mathcal{C}_x(\Sigma_2)$ of reduced cycles, which by Lemma 6.1 is a tree. Lemma 6.7 says that the quotient $\mathcal{C}_x(\Sigma_2)/\mathcal{I}_2$ is also a tree, so by Lemma 6.11 there exists a strict fundamental domain D for the action of $\mathcal{I}_2$ on $\mathcal{C}_x(\Sigma_2)$. Using Corollary 4.2, the edges of $\mathcal{C}_x(\Sigma_2)$ are the nondegenerate cells $\mathcal{X}_x(\gamma)$, where γ is a multicurve that separates Σ_2 into two subsurfaces. Such multicurves have three components, so by Lemma 6.6 the stabilizers in $\mathcal{I}_2$ of the edges in $\mathcal{C}_x(\Sigma_2)$ are trivial (here we are using the *uniqueness* claimed in Lemma 6.6

– the only element that fixes our edge is the identity). Theorem 6.12 therefore says that

$$\mathcal{I}_2 = \mathop{\ast}_{c \in D^{(0)}} (\mathcal{I}_2)_c. \tag{18}$$

Define

$$C = \bigcup_{c \in D^{(0)}} C_c,$$

where C_c is the set of separating curves given by Lemma 6.16. Lemma 6.16 combined with (18) implies that $\mathcal{I}_2$ is a free group with free basis $\{T_\delta \mid \delta \in C\}$. Moreover, combining Lemma 6.16 with Lemma 6.15 we obtain a bijection $\phi : C \to \mathcal{S}$ such that for all $\delta \in C$, the separating twist T_δ induces the separating splitting $\phi(\delta)$, as desired. □

6.4. Curve stabilizers

This section is devoted to the proof of Lemma 6.16. We divide it into two cases which we treat separately.

Lemma 6.17. *Let $x \in \mathrm{H}_1(\Sigma_2; \mathbb{Z})$ be a primitive element and let γ be an oriented simple closed curve on Σ_2 such that $[\gamma] = x$. Then there exists a set C_γ of simple closed separating curves on Σ_2 with the following three properties.*

- *Each $\delta \in C_\gamma$ is disjoint from γ.*
- *The stabilizer subgroup $(\mathcal{I}_2)_\gamma$ is a free group on the free basis $\{T_\delta \mid \delta \in C_\gamma\}$.*
- *There exists a bijection $\phi : C_\gamma \to \mathcal{S}_x$ such that for $\delta \in C_\gamma$, the curve δ induces the separating splitting $\phi(\delta)$.*

Proof. We will use the notation and results from Chapter 5 concerning the stabilizers in Torelli of oriented nonseparating simple closed curves. Recall that $\Sigma_{2,\gamma}$ is the surface that results from cutting Σ_2 open along γ. Let $\{\partial_1, \partial_2\}$ be the boundary components of $\Sigma_{2,\gamma}$ (see Figure 6.6) and let $\psi : \mathcal{I}_{2,\gamma} \to (\mathcal{I}_2)_\gamma$ be the homomorphism obtained by gluing ∂_1 and ∂_2 back together. Lemma 5.10 says that there is a short exact sequence

$$1 \longrightarrow \mathbb{Z} \longrightarrow \mathcal{I}_{2,\gamma} \xrightarrow{\psi} (\mathcal{I}_2)_\gamma \longrightarrow 1,$$

where the kernel $\mathbb{Z}$ is generated by $T_{\partial_1} T_{\partial_2}^{-1}$. Next, let S be the γ-splitting surface depicted in Figure 6.6, let $\widehat{\Sigma}_{2,\gamma}$ be the result of gluing a disc to $\Sigma_{2,\gamma}$ along ∂_1, and let $\pi = \pi_1(\widehat{\Sigma}_{2,\gamma})$. Theorem 5.11 says that there is a decomposition

$$\mathcal{I}_{2,\gamma} = [\pi, \pi] \rtimes \mathcal{I}(S).$$

Since S is a genus 1 surface with 1 boundary component, Lemma 3.3 says that $\mathcal{I}(S) \cong \mathbb{Z}$ with generator the Dehn twist about boundary component of S. Below in Claim 1 we will prove that $\psi : \mathcal{I}_{2,\gamma} \to (\mathcal{I}_2)_\gamma$ restricts to an isomorphism between $[\pi, \pi]$ and $(\mathcal{I}_2)_\gamma$, so most of this proof will concern $[\pi, \pi]$.

Let $\alpha, \beta \in \pi$ be the curves depicted in Figure 6.7. Thus π is a rank 2 free group on α and β. The curve $[\alpha, \beta] \in [\pi, \pi]$ is as shown in Figure 6.7. Using the recipe discussed at the end of §5.3, we see that the element of $\mathcal{I}_{2,\gamma}$ associated to $[\alpha, \beta]$ is $T_{\widehat{\eta}} T_{\partial_2}^{-1} T_{\partial_1}$, where $\widehat{\eta}$ is the boundary component of S (again, see Figure 6.7). It follows that $\psi([\alpha, \beta]) = T_\eta$, where η is the simple closed separating curve depicted in Figure 6.7. We now prove the following.

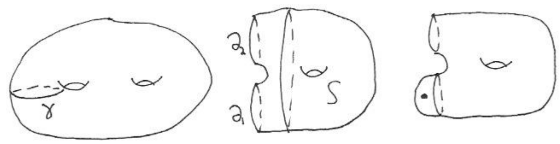

FIGURE 6.6. On the left is an oriented simple closed nonseparating curve γ on Σ_2. In the middle is the cut-open surface $\Sigma_{2,\gamma}$. The boundary components are $\{\partial_1, \partial_2\}$ and S is a γ-splitting surface. On the right is the surface $\widehat{\Sigma}_{2,\gamma}$ that results from gluing a disc to ∂_1. The basepoint for $\pi_1(\widehat{\Sigma}_{2,\gamma})$ in this glued-on disc is as indicated.

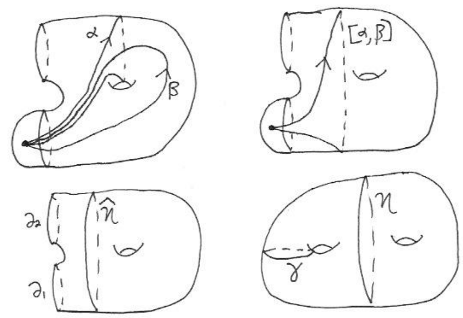

FIGURE 6.7. In the upper-left, generators α and β for $\pi = \pi_1(\widehat{\Sigma}_{2,\gamma})$ are drawn. Their commutator $[\alpha, \beta] = \alpha\beta\alpha^{-1}\beta^{-1}$ is the curve in the upper-right. The element of $\mathcal{I}_{2,\gamma}$ associated to $[\alpha, \beta]$ is $T_{\widehat{\eta}} T_{\partial_2}^{-1} T_{\partial_1}$, where $\widehat{\eta}$ is the curve shown in the lower-left. The key point here is that $\widehat{\eta}$ and ∂_2 are the boundary components of a regular neighborhood of $[\alpha, \beta]$ (see the recipe for this at the end of §5.3). Since ψ is induced by the map that glues ∂_1 and ∂_2 back together, we have $\psi(T_{\partial_2}^{-1} T_{\partial_1}) = 1$, and thus $\psi([\alpha, \beta]) = T_\eta$, where η is the curve shown in the lower-right.

CLAIM 1. *The restriction of the surjection* $\psi : \mathcal{I}_{g,\gamma} \to (\mathcal{I}_g)_\gamma$ *to* $[\pi, \pi] < \mathcal{I}_{g,\gamma}$ *is an isomorphism.*

PROOF OF CLAIM. As was discussed above, the kernel of ψ is $\mathbb{Z}$ with generator $T_{\partial_1} T_{\partial_2}^{-1}$. It is clear that no nontrivial power of $T_{\partial_1} T_{\partial_2}^{-1}$ lies in $[\pi, \pi] < \mathcal{I}_{g,\gamma}$; indeed, this follows immediately from the fact that $[\pi, \pi]$ consists of all mapping classes in $\mathcal{I}_{g,\gamma}$ that become trivial when a disc is glued to $\Sigma_{g,\gamma}$ along ∂_1. We deduce that ψ restricts to an injection $[\pi, \pi] \hookrightarrow (\mathcal{I}_g)_\gamma$. To see that this injection is a surjection, since $\mathcal{I}_{g,\gamma} = [\pi, \pi] \rtimes \mathcal{I}(S)$ it is enough to show that its image contains $\psi(\mathcal{I}(S))$. As we discussed above, $\mathcal{I}(S) \cong \mathbb{Z}$ with generator the Dehn twist about the boundary component of S. The map ψ takes this Dehn twist to T_η, which is also $\psi([\alpha, \beta])$. The claim follows. □

We now focus on $[\pi, \pi]$. If G is a group and $x, y \in G$, then let x^y denote yxy^{-1}.

CLAIM 1. *The group* $[\pi, \pi]$ *is a free group with free basis* $\{[\alpha, \beta]^{\alpha^k \beta^\ell} \mid k, \ell \in \mathbb{Z}\}$.

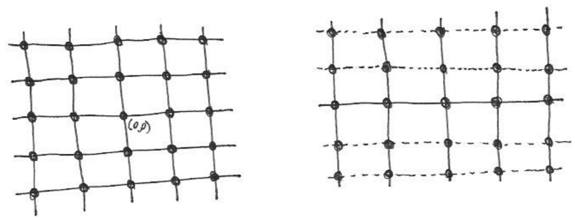

FIGURE 6.8. On the left is the cover X of the wedge of two circles corresponding to $[\pi, \pi]$. The horizontal edges map to the circle corresponding to α (going from left to right) and the vertical edges map to the circle corresponding to β (going from bottom to top). On the right is a maximal tree T in X; the omitted edges are dashed. Identifying X with a subset of $\mathbb{R}^2$, the tree T consists of the horizontal line $y = 0$ and the vertical lines $x = n$ for all $n \in \mathbb{Z}$.

PROOF OF CLAIM. The group π is a rank two free group with free basis $\{\alpha, \beta\}$. This is the fundamental group of a wedge of two circles. The cover of this wedge of two circles corresponding to $[\pi, \pi]$ is the "grid" X shown in Figure 6.8. In words, X is a graph with vertex set $\mathbb{Z}^2$ and with edges connecting (n, m) to $(n+1, m)$ and $(n, m+1)$ for all $(n, m) \in \mathbb{Z}^2$. Let $T \subset X$ be the maximal tree depicted in Figure 6.8. For $(i, j) \in \mathbb{Z}^2$, let $a_{i,j}$ be the edge in X connecting (i, j) and $(i+1, j)$ and let $s_{i,j} \in [\pi, \pi]$ be the element corresponding to the loop in X that starts at $(0, 0)$, goes along the unique path in T to (i, j), then goes along a_{ij} to $(i+1, j)$, and then goes along the unique path in T connecting $(i+1, j)$ to $(0, 0)$. The edges of X that do not lie in T are

$$\{a_{i,j} \mid (i, j) \in \mathbb{Z}^2,\ j \neq 0\},$$

so $[\pi, \pi]$ is a free group with free basis

$$\{s_{i,j} \mid (i, j) \in \mathbb{Z}^2,\ j \neq 0\}.$$

With respect to this basis, for $(k, \ell) \in \mathbb{Z}^2$ we have

$$[\alpha, \beta]^{\alpha^k \beta^\ell} = \alpha^k \beta^\ell \alpha \beta \alpha^{-1} \beta^{-1} \beta^{-\ell} \alpha^{-k} = \begin{cases} s_{k,\ell} s_{k,\ell+1}^{-1} & \text{if } \ell \neq 0, \\ s_{k,\ell+1}^{-1} & \text{if } \ell = 0. \end{cases}$$

It follows immediately that the desired set is also a free basis for $[\pi, \pi]$. □

We now determine the image under ψ of the generators for $[\pi, \pi]$ given by Claim 2.

CLAIM 2. *For $k, \ell \in \mathbb{Z}$, we have $\psi([\alpha, \beta]^{\alpha^k \beta^\ell}) = T_{\eta_{k,\ell}}$, where $\eta_{k,\ell}$ is a separating curve inducing the separating splitting*

$$(\langle a_1, b_1 + k a_2 + \ell b_2 \rangle, \langle a_2 + \ell a_1, b_2 - k a_1 \rangle).$$

PROOF OF CLAIM. Recall that $[\pi, \pi]$ is embedded in $\mathcal{I}_{2,\gamma}$ as follows. The group $\mathrm{Mod}_{2,\gamma}$ contains the "disc-pushing" subgroup, which is isomorphic to $\pi_1(U\widehat{\Sigma}_{2,\gamma})$. Here $U\widehat{\Sigma}_{2,\gamma}$ is the unit tangent bundle of $\widehat{\Sigma}_{2,\gamma}$. There is a natural projection map $\pi_1(U\widehat{\Sigma}_{2,\gamma}) \to \pi_1(\widehat{\Sigma}_{2,\gamma}) = \pi$. For every curve $\zeta \in [\pi, \pi]$, there is a unique $\widetilde{\zeta}$ in the disc-pushing subgroup $\pi_1(U\widehat{\Sigma}_{2,\gamma})$ of $\mathrm{Mod}_{2,\gamma}$ that projects to ζ such that $\widetilde{\zeta} \in \mathcal{I}_{2,\gamma}$; this is the element of $\mathcal{I}_{2,\gamma}$ corresponding to ζ.

As we said above, the element of $\mathcal{I}_{2,\gamma}$ corresponding to $[\alpha, \beta]$ is $T_{\widehat{\eta}} T_{\partial_2}^{-1} T_{\partial_1}$. Now let $\theta_{k,\ell}$ be an arbitrary element of the disc-pushing subgroup $\pi_1(U\widehat{\Sigma}_{2,\gamma})$ that projects to $\alpha^k \beta^\ell \in \pi$. Since $\mathcal{I}_{2,\gamma}$ is a normal subgroup of $\mathrm{Mod}_{2,\gamma}$, we have that

$$\theta_{k,\ell}(T_{\widehat{\eta}} T_{\partial_2}^{-1} T_{\partial_1})\theta_{k,\ell}^{-1} \in \mathcal{I}_{2,\gamma}.$$

Thus $\theta_{k,\ell}(T_{\widehat{\eta}} T_{\partial_2}^{-1} T_{\partial_1})\theta_{k,\ell}^{-1}$ is the element of $\mathcal{I}_{2,\gamma}$ corresponding to $[\alpha, \beta]^{\alpha^k \beta^\ell}$. It follows that

$$\psi([\alpha, \beta]^{\alpha^k \beta^\ell}) = \psi(\theta_{k,\ell}) T_\eta \psi(\theta_{k,\ell})^{-1} = T_{\psi(\theta_{k,\ell})(\eta)}.$$

We thus can take $\eta_{k,\ell} = \psi(\theta_{k,\ell})(\eta)$. Since η is a separating curve inducing the separating splitting $(\langle a_1, b_1 \rangle, \langle a_2, b_2 \rangle)$, we deduce that $\eta_{k,\ell}$ is a separating curve inducing the separating splitting

$$(\langle (\theta_{k,\ell})_*(a_1), (\theta_{k,\ell})_*(b_1) \rangle, \langle (\theta_{k,\ell})_*(a_2), (\theta_{k,\ell})_*(b_2) \rangle).$$

Using Lemma 5.8, we see that this is

$$(\langle a_1, b_1 + k a_2 + \ell b_2 \rangle, \langle a_2 + \ell a_1, b_2 - k a_1 \rangle),$$

as desired. □

Define $C_\gamma = \{\eta_{k,\ell} \mid k, \ell \in \mathbb{Z}\}$. Combining the three claims above, we see that $(\mathcal{I}_2)_\gamma$ is the free group on the set $\{T_\delta \mid \delta \in C_\gamma\}$. Next, define $\phi : C_\gamma \to \mathcal{S}_x$ via the formula

$$\phi(\eta_{k,\ell}) = (\langle a_1, b_1 + k a_2 + \ell b_2 \rangle, \langle a_2 + \ell a_1, b_2 - k a_1 \rangle).$$

Thus for $\delta \in C_\gamma$, the separating splitting induced by δ is $\phi(\delta)$. We then have the following.

Claim 3. *The map ϕ is a bijection.*

Proof of claim. It is enough to show that ϕ is surjective. Consider $(U_1, U_2) \in \mathcal{S}_x$, so $a_1 = x \in U_1$. We must have $U_1 = \langle a_1, w \rangle$, where $w \in \mathrm{H}_1(\Sigma_2; \mathbb{Z})$ satisfies $\hat{i}(a_1, w) = 1$. Expanding w out in terms of our symplectic basis $\{a_1, b_1, a_2, b_2\}$ for $\mathrm{H}_1(\Sigma_2; \mathbb{Z})$, the fact that $\hat{i}(a_1, w) = 1$ implies that $w = b_1 + k a_2 + \ell b_2$ for some $k, \ell \in \mathbb{Z}$. Also, using the fact that

$$U_2 = \{v \in \mathrm{H}_1(\Sigma_2; \mathbb{Z}) \mid \hat{i}(a_1, v) = \hat{i}(w, v) = 0\},$$

we see that $U_2 = \langle a_2 + \ell a_1, b_2 - k a_1 \rangle$. Thus $(U_1, U_2) = \phi(\eta_{k,\ell})$, as desired. □

This completes the proof of Lemma 6.17. □

Lemma 6.18. *Let $x \in \mathrm{H}_1(\Sigma_2; \mathbb{Z})$ be a primitive element and let $c_1\gamma_1 + c_2\gamma_2$ be a vertex of $\mathcal{C}_x(\Sigma_2)$ such that $c_1, c_2 \neq 0$. Then there exists a set $C_{c_1\gamma_1 + c_2\gamma_2}$ of simple closed separating curves on Σ_2 with the following three properties.*

- *Each $\delta \in C_{c_1\gamma_1 + c_2\gamma_2}$ is disjoint from $\gamma_1 \cup \gamma_2$.*
- *The stabilizer subgroup $(\mathcal{I}_2)_{c_1\gamma_1 + c_2\gamma_2}$ is a free group on the free basis*

$$\{T_\delta \mid \delta \in C_{c_1\gamma_1 + c_2\gamma_2}\}.$$

- *Let $c_1[\![v_1]\!] + c_2[\![v_2]\!]$ be the image of $c_1\gamma_1 + c_2\gamma_2$ in $\mathcal{C}_x(\Sigma_2)/\mathcal{I}_2$. Then there exists a bijection $\phi : C_{c_1\gamma_1 + c_2\gamma_2} \to \mathcal{S}_{c_1[\![v_1]\!] + c_2[\![v_2]\!]}$ such that for $\delta \in C_{c_1\gamma_1 + c_2\gamma_2}$, the curve δ induces the separating splitting $\phi(\delta)$.*

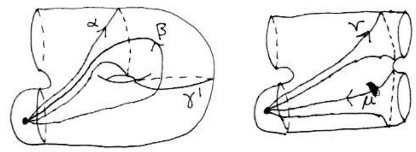

FIGURE 6.9. On the left is the surface $\widehat{\Sigma}_{2,\gamma}$ with the curve γ' on it together with the elements $\alpha, \beta \in \pi$. On the right is the surface X obtained by cutting $\widehat{\Sigma}_{2,\gamma}$ open along γ'. The curves $\nu, \mu \in \pi_1(X)$ form a free basis for the rank 2 free group $\pi_1(X)$, and the natural map $\pi_1(X) \to \pi$ that glues the boundary components back together takes ν to α and $\nu\mu^{-1}$ to $[\alpha, \beta]$.

PROOF. Clearly we have

$$(\mathcal{I}_2)_{c_1\gamma_1+c_2\gamma_2} = (\mathcal{I}_2)_{\gamma_1\cup\gamma_2} = ((\mathcal{I}_2)_{\gamma_1})_{\gamma_2}.$$

To identify this with a subgroup of the group described in Lemma 6.17, we define $\gamma = \gamma_1$. Let the notation be as in the proof of Lemma 6.17, so $(\mathcal{I}_2)_\gamma \cong [\pi, \pi]$, where $\pi = \pi_1(\widehat{\Sigma}_{2,\gamma})$. Let γ' be the curve in $\Sigma_{2,\gamma} \subset \widehat{\Sigma}_{2,\gamma}$ that maps to γ_2 under the map $\Sigma_{2,\gamma} \to \Sigma_2$ that glues the boundary components ∂_1 and ∂_2 back together. We then have

$$((\mathcal{I}_2)_\gamma)_{\gamma_2} \cong \{\zeta \in [\pi, \pi] \mid \zeta \text{ is disjoint from } \gamma'\}.$$

To simplify our notation, let

$$\Gamma = \{\zeta \in [\pi, \pi] \mid \zeta \text{ is disjoint from } \gamma'\}.$$

Assume that the curves α and β from the proof of Lemma 6.17 are chosen as in Figure 6.9. Also, choose the expression $(\langle a_1, b_1\rangle, \langle a_2, b_2\rangle)$ for the separating splitting induced by η such that $a_2 = [\gamma_2] = v_2$ (recall that $a_1 = [\gamma] = v_1$). It is then enough to prove two things. The first is as follows.

CLAIM 1. *The group Γ is a free group with free basis $\{[\alpha, \beta]^{\alpha^k} \mid k \in \mathbb{Z}\}$.*

PROOF OF CLAIM. Let X be the result of cutting $\widehat{\Sigma}_{2,\gamma}$ open along γ' and let $\nu, \mu \in \pi_1(X)$ be the curves shown in Figure 6.9. The natural map $\pi_1(X) \to \pi$ takes ν to α and $\nu\mu^{-1}$ to $[\alpha, \beta]$. The group $\pi_1(X)$ is a rank 2 free group with basis $\{\nu, \mu\}$. Also, letting Γ' be the kernel of the map $\pi_1(X) \to \mathbb{Z}$ that takes ν and μ both to 1, it is clear that the map $\pi_1(X) \to \pi$ takes Γ' isomorphically onto Γ. Regarding $\pi_1(X)$ as the fundamental group of a wedge of two circles, the cover corresponding to Γ' is as shown in Figure 6.10. From this cover, it is clear that Γ' is a free group with free basis

$$\{\nu^k\mu\nu^{-(k+1)} \mid k \in \mathbb{Z}\} = \{\nu^k(\nu^{-1}\mu)\nu^{-k} \mid k \in \mathbb{Z}\}.$$

This free basis maps to the free basis

$$\{([\alpha, \beta]^{-1})^{\alpha^k} \mid k \in \mathbb{Z}\}$$

for Γ. The claim follows. □

For the second claim that must be proved, recall that $\mathcal{S}_{c_1[\![v]\!]+c_2[\![v']\!]}$ is the set of all separating splittings $(U_1, U_2) \in \mathcal{S}$ such that $a_1 = v \in U_1$ and $a_2 = v' \in U_2$. We then have the following.

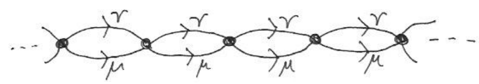

FIGURE 6.10. The cover of the wedge of two circles whose fundamental group is Γ'. The edges of this graph map to the circles corresponding to ν and μ as indicated.

CLAIM 2. *The set* $\mathcal{S}_{c_1[\![v]\!]+c_2[\![v']\!]}$ *of separating splittings equals*

$$\{(\langle a_1, b_1 + ka_2\rangle, \langle a_2, b_2 - ka_1\rangle) \mid k \in \mathbb{Z}\}.$$

PROOF OF CLAIM. An arbitrary element of $\mathcal{S}_{c_1[\![v]\!]+c_2[\![v']\!]}$ is of the form

$$(\langle a_1, w_1\rangle, \langle a_2, w_2\rangle)$$

for some $w_1, w_2 \in \mathrm{H}_1(\Sigma_2; \mathbb{Z})$ such that

$$\hat{i}(a_1, w_1) = \hat{i}(a_2, w_2) = 1 \quad \text{and} \quad \hat{i}(a_1, w_2) = \hat{i}(a_2, w_1) = \hat{i}(w_1, w_2) = 0.$$

Since $\hat{i}(a_1, w_1) = 1$ and $\hat{i}(a_2, w_1) = 0$, we have $w_1 = b_1 + \ell a_1 + ka_2$ for some $\ell, k \in \mathbb{Z}$. Since all we care about is the span of a_1 and w_1, we can assume that $\ell = 0$. Similarly, we have $w_2 = b_2 + k'a_1$ for some $k' \in \mathbb{Z}$. Finally, since $\hat{i}(w_1, w_2) = 0$ we must have $k' = -k$, as desired. □

Lemma 6.18 follows immediately. □

Part 3

The Johnson Homomorphism

CHAPTER 9

A commutator quotient

In this chapter, we prove the following theorem, which will play an important role in the construction of the Johnson homomorphism. Recall that F_n is the free group on n letters. For $x \in F_n$, denote by $[x]$ the associated element of $F_n^{\text{ab}} = \mathbb{Z}^n$.

THEOREM 9.1. *There exists a homomorphism $\rho : [F_n, F_n] \to \wedge^2 \mathbb{Z}^n$ such that*

$$\rho([x, y]) = [x] \wedge [y] \qquad (x, y \in F_n).$$

We will give three proofs of Theorem 9.1. The first is in §9.1 and uses Magnus expansions, which are certain homomorphisms from a free group to a truncation of a tensor algebra. The second in in §9.2 and uses the Fox free differential calculus, which will also be used later when we study cup products. The third proof is in §9.3 and uses some tools from group cohomology.

Thoughout this section, we will give a more complete exposition of the tools we use than is strictly necessary for the proof of Theorem 9.1. These tools are important in many contexts, and a geodesic path to Theorem 9.1 would not necessarily be the most enlightening one.

Since we will need them later, we pause now to record some naturality properties of the homomorphism ρ from Theorem 9.1. Observe that $\text{Aut}(F_n)$ acts on $\wedge^2 \mathbb{Z}^n$ via its action on $F_n^{\text{ab}} = \mathbb{Z}^n$.

LEMMA 9.2. *For $w \in [F_n, F_n]$ and $f \in \text{Aut}(F_n)$, we have $\rho(f(w)) = f(\rho(w))$.*

PROOF. Clearly it is enough to prove this for $w = [x, y]$ with $x, y \in F_n$. But then

$$\rho(f(w)) = \rho([f(x), f(y)]) = [f(x)] \wedge [f(y)] = f([x] \wedge [y]) = f(\rho(w)). \qquad \square$$

This has the following corollary.

COROLLARY 9.3. *Assume that $f \in \text{Aut}(F_n)$ acts trivially on $\mathbb{Z}^n$. Then*

$$\rho(f(w)) = \rho(w) \qquad (w \in [F_n, F_n]).$$

This holds in particular for inner automorphisms, so for all $x \in F_n$ we have

$$\rho(xwx^{-1}) = \rho(w) \qquad (w \in [F_n, F_n]).$$

9.1. Via Magnus expansions

In this section, we introduce Magnus expansions and use them to prove Theorem 9.1. Magnus introduced these homomorphisms in [**Mag35**] and used them to prove that the intersection of the lower central series of a free group is trivial (see Theorem 9.12 below). See [**MKS76**, Chapter 5] for more details concerning Magnus expansions.

The truncated tensor algebra Let $T(\mathbb{Z}^n)$ be the tensor algebra of $\mathbb{Z}^n$, so

$$T(\mathbb{Z}^n) = \bigoplus_{i=0}^{\infty} T^i(\mathbb{Z}^n) \quad \text{with } T^i(\mathbb{Z}^n) = (\mathbb{Z}^n)^{\otimes i}.$$

The algebra structure on $T(\mathbb{Z}^n)$ comes from the tensor product. For all $k \geqslant 0$, the algebra $T(\mathbb{Z}^n)$ contains an ideal

$$I_k(\mathbb{Z}^n) = \bigoplus_{i=k}^{\infty} T^i(\mathbb{Z}^n).$$

The *degree k truncated tensor algebra* of $\mathbb{Z}^n$, denoted $\mathcal{A}_k(\mathbb{Z}^n)$, is $T(\mathbb{Z}^n)/I_{k+1}(\mathbb{Z}^n)$. We will regard $\mathcal{A}_k(\mathbb{Z}^n)$ as composed of expressions

$$f = f_0 + f_1 + \cdots + f_k \qquad (f_i \in T^i(\mathbb{Z}^n))$$

and will call f_i the *degree i component* of f. Observe that

$$(f_0 + f_1 + \cdots + f_k)(g_0 + g_1 + \cdots g_k) = \sum_{i+j \leqslant k} f_i g_j.$$

Invertible elements. The following lemma gives many invertible elements in $\mathcal{A}_k(\mathbb{Z}^n)$.

LEMMA 9.4. *If the degree 0 component of $f \in \mathcal{A}_k$ is 1, then f is invertible. Moreover, the degree 0 component of f^{-1} is also 1.*

PROOF. Write $f = 1 + f_1 + \cdots + f_k$. Our goal is to find some $g = g_0 + g_1 + \cdots + g_k \in \mathcal{A}_k$ such that

$$(1 + f_1 + \cdots + f_k)(g_0 + g_1 + \cdots + g_k) = 1. \tag{19}$$

We can solve for the g_j inductively as follows. First, $g_0 = 1$. Second, if we have already found $g_0, \ldots, g_{j-1}$, then (19) implies that

$$g_j + f_1 g_{j-1} + f_2 g_{j-2} + \cdots + f_j g_0 = 0,$$

so

$$g_j = -f_j g_0 - f_{j-1} g_1 - \cdots - f_1 g_{j-1}.$$ □

The Magnus expansion. Define $\mathcal{A}'_k(\mathbb{Z}^n)$ to be the set of all elements of $\mathcal{A}_k(\mathbb{Z}^n)$ whose degree 0 components are 1. Lemma 9.4 says that $\mathcal{A}'_k(\mathbb{Z}^n)$ forms a group under multiplication. Letting $\{x_1, \ldots, x_n\}$ be a fixed free basis for F_n, there thus exists a homomorphism

$$\psi_k : F_n \to \mathcal{A}'_k(\mathbb{Z}^n)$$

taking $x_i \in F_n$ to $1 + [x_i] \in \mathcal{A}_k(\mathbb{Z}^n)$ for $1 \leqslant i \leqslant n$. We will call ψ_k the *degree k Magnus expansion* of F_n.

REMARK 9.5. The Magnus expansion depends on the free basis $\{x_1, \ldots, x_n\}$.

REMARK 9.6. For all $k \geqslant 1$, the homomorphism $\psi_k : F_n \to \mathcal{A}'_k(\mathbb{Z}^n)$ is the composition of $\psi_{k+1} : F_n \to \mathcal{A}'_{k+1}(\mathbb{Z}^n)$ with the natural projection $\mathcal{A}'_{k+1}(\mathbb{Z}^n) \to \mathcal{A}'_k(\mathbb{Z}^n)$. Letting $\mathcal{A}'_\infty(\mathbb{Z}^n)$ be the inverse limit of the $\mathcal{A}'_k(\mathbb{Z}^n)$, the ψ_k thus piece together to yield a homomorphism $\psi_\infty : F_n \to \mathcal{A}'_\infty(\mathbb{Z}^n)$. Most authors only call ψ_∞ a Magnus expansion, but we find it more straightforward to work with the individual ψ_k.

Abelian quotients on kernels. The following lemma shows that the degree 1 Magnus expansion $\psi_1 : F_n \to \mathcal{A}'_1(\mathbb{Z}^n)$ can be identified with the abelianization map $F_n \to \mathbb{Z}^n$.

LEMMA 9.7. *For all $x \in F_n$ we have $\psi_1(x) = 1 + [x] \in \mathcal{A}'_1(\mathbb{Z}^n)$. Consequently, we have* $\ker(\psi_1) = [F_n, F_n]$.

PROOF. Using the recipe in the proof of Lemma 9.4 for computing inverses in $\mathcal{A}_k$, we see that in $\mathcal{A}'_1(\mathbb{Z}^n)$ we have

$$(1 + [x_i])^{-1} = 1 - [x_i] \qquad (1 \leqslant i \leqslant n).$$

Therefore, given a word $w = x_{i_1}^{\epsilon_1} \cdots x_{i_\ell}^{\epsilon_\ell}$ in F_n with $1 \leqslant i_j \leqslant n$ and $\epsilon_j \in \{1, -1\}$ for $1 \leqslant j \leqslant \ell$, we have

$$\psi_1(w) = (1 + \epsilon_1[x_{i_1}]) \cdots (1 + \epsilon_\ell[x_{i_\ell}]) = 1 + (\epsilon_1[x_{i_1}] + \cdots + \epsilon_\ell[x_{i_\ell}]) = 1 + [w],$$

as desired. Observe that in the above calculation we discarded all components of degree 2 and higher. □

The following lemma generalizes the main insight of Lemma 9.7.

LEMMA 9.8. *Fix $k \geqslant 2$. Then there exists a homomorphism $\phi_k : \ker(\psi_{k-1}) \to T^k(\mathbb{Z}^n)$ such that*

$$\psi_k(w) = 1 + \phi_k(w) \qquad (w \in \ker(\psi_{k-1})).$$

PROOF. Let $\pi : \mathcal{A}'_k(\mathbb{Z}^n) \to \mathcal{A}'_{k-1}(\mathbb{Z}^n)$ be the natural projection homomorphism, so $\psi_{k-1} = \pi \circ \psi_k$. Since

$$\ker(\pi) = \{1 + f_k \mid f_k \in T^k(\mathbb{Z}^n)\},$$

there exists a set map $\phi_k : \ker(\psi_{k-1}) \to T^k(\mathbb{Z}^n)$ such that

$$\psi_k(w) = 1 + \phi_k(w) \qquad (w \in \ker(\psi_{k-1})).$$

The fact that ϕ_k is a homomorphism follows from the easy calculation

$$(1 + f_k)(1 + g_k) = 1 + (f_k + g_k) \qquad (f_k, g_k \in T^k(\mathbb{Z}^n))$$

in $\mathcal{A}'_k(\mathbb{Z}^n)$; here we discard all components of degree $k+1$ and higher. □

Deeper in the lower central series. We will not need the results in this paragraph later in the book, but it would be strange to not include them in a discussion of Magnus expansions. We start with the following lemma.

LEMMA 9.9. *Fix $k \geqslant 2$. Then for all $x \in F_n$ and $w \in \ker(\psi_{k-1})$ we have $\psi_k(xwx^{-1}) = \psi_k(w)$.*

PROOF. Using Lemma 9.8, we have $\psi_k(w) = 1 + f_k$ for some $f_k \in T^k(\mathbb{Z}^n)$. Write

$$\psi_k(x) = 1 + g_1 + \cdots + g_k \qquad (g_i \in T^i(\mathbb{Z}^n)).$$

We then have

$$\psi_k(x) f_k = f_k + g_1 f_k + \cdots + g_k f_k = f_k.$$

Similarly, we have $f_k \psi_k(x^{-1}) = f_k$. Therefore,

$$\psi_k(xwx^{-1}) = \psi_k(x)(1 + f_k)\psi_k(x^{-1}) = (\psi_k(x) + f_k)\psi_k(x^{-1}) = 1 + f_k = \psi_k(w),$$

as desired. □

This has the following corollary. The *lower central series* of a group G is the sequence of subgroups defined inductively by

$$\gamma_1(G) = G \quad \text{and} \quad \gamma_{k+1}(G) = [G, \gamma_k(G)].$$

COROLLARY 9.10. *For all $k \geqslant 1$, we have $\gamma_{k+1}(F_n) \subset \ker(\psi_k)$.*

PROOF. The proof will be by induction on k. The base case $k = 1$ is Lemma 9.7. Assume now that $\gamma_k(F_n) \subset \ker(\psi_{k-1})$. Consider $x \in F_n$ and $w \in \gamma_k(F_n)$. We must show that $[x, w] \in \ker(\psi_k)$. By Lemma 9.9 we have

$$\psi_k([x, w]) = \psi_k(xwx^{-1})\psi_k(w^{-1}) = \psi_k(w)\psi_k(w^{-1}) = 1,$$

as desired. □

REMARK 9.11. Witt [**Wit37**] proved that for all $k \geqslant 1$, the group $\ker(\psi_k)$ is actually equal to $\gamma_{k+1}(F_n)$. In addition to the original source, see [**MKS76**, Chapter 5] and [**Ser92**, Chapter IV.6] and [**CFL58**] for the details of this.

We finally deduce the following theorem of Magnus [**Mag35**]. For an alternate topological proof, see [**MP10**].

THEOREM 9.12. *The group F_n is* residually nilpotent*, that is, $\cap_{k=1}^{\infty}\gamma_k(F_n) = 1$.*

PROOF. Consider $w \in F_n$ such that $w \neq 1$. Write $w = x_{i_1}^{m_1} \cdots x_{i_k}^{m_k}$ for some $1 \leqslant i_j \leqslant n$ and $m_j \in \mathbb{Z}\setminus\{0\}$ satisfying $i_j \neq i_{j+1}$ for $1 \leqslant j < k$. For $1 \leqslant j \leqslant k$, observe that

$$\psi_k(x_{i_j}^{m_j}) = (1 + [x_{i_j}])^{m_j}$$

is 1 plus a $\mathbb{Z}$-linear combination of terms of the form $[x_{i_j}]^p$ with $1 \leqslant p \leqslant k$. Also, the coefficient of $[x_{i_j}]$ is m_j. Consider the $T^k(\mathbb{Z}^n)$ term of

$$\psi_k(w) = \psi_k(x_{i_1}^{m_1}) \cdots \psi_k(x_{i_k}^{m_k}).$$

Expressing this in terms of the basis

$$\{[x_{p_1}][x_{p_2}] \cdots [x_{p_k}] \mid 1 \leqslant p_j \leqslant n \text{ for } 1 \leqslant j \leqslant k\},$$

the only term that does not involve a basis element with $p_j = p_{j+1}$ for some $1 \leqslant j < k$ is

$$m_1 m_2 \cdots m_k [x_{i_1}][x_{i_2}] \cdots [x_{i_k}].$$

This implies in particular that $\psi_k(w) \neq 1$, so by Corollary 9.10 we have $w \notin \gamma_{k+1}(w)$. □

The quotient of the commutator subgroup. We finally prove Theorem 9.1.

PROOF OF THEOREM 9.1. Combining Lemmas 9.7 and 9.8, we obtain a homomorphism

$$\rho : [F_n, F_n] \to T^2(\mathbb{Z}^n)$$

such that $\psi_2(w) = 1 + \rho(w)$ for all $w \in [F_n, F_n]$. We will prove that the image of ρ lies in $\wedge^2\mathbb{Z}^n \subset T^2(\mathbb{Z}^n)$. To do this, it is enough to prove the formula claimed in the theorem, namely $\rho([x, y]) = [x] \wedge [y] \in \wedge^2\mathbb{Z}^n$ for all $x, y \in F_n$. Write

$$\psi_2(x) = 1 + [x] + f_2 \quad \text{and} \quad \psi_2(y) = 1 + [y] + g_2$$

for some $f_2, g_2 \in T^2(\mathbb{Z}^n)$. Using the recipe from the proof of Lemma 9.4, we have

$$(1+[x]+f_2)^{-1} = 1-[x]+([x][x]-f_2) \quad \text{and} \quad (1+[y]+g_2)^{-1} = 1-[y]+([y][y]-g_2).$$

We then have that

$$\begin{aligned}\psi_2([x,y]) &= (1+[x]+f_2)(1+[y]+g_2)(1+[x]+f_2)^{-1}(1+[y]+g_2)^{-1}\\ &= (1+[x]+f_2)(1+[y]+g_2)\\ &\qquad (1-[x]+([x][x]-f_2))(1-[y]+([y][y]-g_2))\\ &= 1+[x][y]-[y][x]\\ &= 1+[x]\wedge[y],\end{aligned}$$

as desired. Here we discard all components of degree 3 and higher. □

9.2. Via the Fox free differential calculus

We next show how to prove Theorem 9.1 via the Fox free differential calculus, which was introduced by Fox in [**Fox53**]. This proof will appear quite different from the proof in §9.1, but in reality it is very similar. We will comment on the connection between the two approaches at the end of this section.

Derivations. Let G be a group and M be a G-module. A *derivation* from G to M is a function $\phi : G \to M$ such that

$$\phi(gh) = \phi(g) + g \cdot \phi(h) \qquad (g, h \in G).$$

Derivations are also sometimes called *crossed homomorphisms*.

EXAMPLE 9.13. If the action of G on M is trivial, then derivations from G to M are the same as homomorphisms.

EXAMPLE 9.14. For $m \in M$, define a function $\phi_m : G \to M$ via the formula $\phi_m(g) = g \cdot m - m$. Then ϕ_m is a derivation; indeed,

$$\phi_m(gh) = gh \cdot m - m = (g \cdot m - m) + (gh \cdot m - g \cdot m) = \phi_m(g) + g \cdot \phi_m(h).$$

The derivation ϕ_m is often called a *principal derivation*.

The following lemma gives a useful alternate definition of a derivation.

LEMMA 9.15. *Let G be a group and M be a G-module. Then a set map $\phi : G \to M$ is a derivation if and only if the map $(\phi, id) : G \to M \rtimes G$ is a homomorphism.*

PROOF. The map (ϕ, id) is a homomorphism if and only if for all $g, h \in G$ we have

$$(\phi(gh), gh) = (\phi(g), g)(\phi(h), h) = (\phi(g) + g \cdot \phi(h), gh).$$

The second equality follows from the definition of a semidirect product. □

This lemma has the following corollary. Let F_n be the free group on the set $\{x_1, \ldots, x_n\}$.

COROLLARY 9.16. *If M is an F_n-module and $m_1, \ldots, m_n \in M$, then there exists a unique derivation $\phi : F_n \to M$ such that $\phi(x_i) = m_i$ for $1 \leqslant i \leqslant n$.*

PROOF. The universal property of a free group says that there exists a unique homomorphism $\psi : F_n \to M \rtimes F_n$ such that $\psi(x_i) = (m_i, x_i)$ for $1 \leqslant i \leqslant n$. The corollary now follows from Lemma 9.15. □

Free derivatives. A *free derivative* on F_n is a derivation $\phi : F_n \to \mathbb{Z}[F_n]$. Denote the set of all free derivatives on F_n by Der_n. Clearly Der_n is closed under addition. Moreover, if $\phi \in \mathrm{Der}_n$ and $\tau \in \mathbb{Z}[F_n]$, then the map $\phi^{(\tau)} : F_n \to \mathbb{Z}[F_n]$ defined by the formula

$$\phi^{(\tau)}(g) = \phi(g) \cdot \tau$$

is a free derivative. Thus Der_n is a right $\mathbb{Z}[F_n]$-module. Its most important elements are the free derivatives $\frac{\partial}{\partial x_i}$ for $1 \leqslant i \leqslant n$ defined via the formula

$$\frac{\partial}{\partial x_i}(x_j) = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{otherwise} \end{cases}$$

for $1 \leqslant j \leqslant n$. Corollary 9.16 says that these exist and are uniquely defined by the above formula. We then have the following.

LEMMA 9.17. *The module Der_n is a free right $\mathbb{Z}[F_n]$-module on the basis $\{\frac{\partial}{\partial x_i} \mid 1 \leqslant i \leqslant n\}$.*

PROOF. Consider $\phi \in \mathrm{Der}_n$. For $1 \leqslant i \leqslant n$, set $\tau_i = \phi(x_i)$. Define

$$\phi' = (\frac{\partial}{\partial x_1})^{(\tau_1)} + \cdots (\frac{\partial}{\partial x_n})^{(\tau_n)}.$$

Then $\phi'(x_i) = \phi(x_i)$ for $1 \leqslant i \leqslant n$. Corollary 9.16 therefore says that $\phi' = \phi$. Morover, it is clear that ϕ' is the only $\mathbb{Z}[F_n]$-linear combination of the $\frac{\partial}{\partial x_i}$ with this property. □

Basic properties of free derivatives. The following lemma summarizes some basic facts about free derivatives.

LEMMA 9.18. *Consider $\phi \in Der_n$.*

(1) *We have $\phi(1) = 0$.*
(2) *For $x \in F_n$, we have $\phi(x^{-1}) = -x^{-1}\phi(x)$. In particular, for $1 \leqslant i, j \leqslant n$ we have*

$$\frac{\partial}{\partial x_i}(x_j^{-1}) = \begin{cases} -x_j^{-1} & \text{if } j = i, \\ 0 & \text{otherwise.} \end{cases}$$

PROOF. For the first claim, observe that

$$\phi(1) = \phi(1 \cdot 1) = \phi(1) + 1 \cdot \phi(1) = 2\phi(1),$$

so $\phi(1) = 0$. For the second claim, observe that

$$0 = \phi(x^{-1}x) = \phi(x^{-1}) + x^{-1}\phi(x),$$

so $\phi(x^{-1}) = -x^{-1}\phi(x)$. □

Augmentations and extending to the group ring. Let $\alpha : \mathbb{Z}[F_n] \to \mathbb{Z}$ be the *augmentation map*, that is, the unique $\mathbb{Z}$-linear map that takes each $g \in F_n$ to 1. We then have the following. If $\phi \in \mathrm{Der}_n$, then observe that ϕ can be linearly extended to a map $\phi : \mathbb{Z}[F_n] \to \mathbb{Z}[F_n]$.

LEMMA 9.19. *If $\phi \in Der_n$ and $\theta, \delta \in \mathbb{Z}[F_n]$, then $\phi(\theta\delta) = \phi(\theta)\alpha(\delta) + \theta\phi(\delta)$.*

PROOF. This holds by definition if $\theta, \delta \in F_n$. To extend this to general elements of the group ring, simply observe that both sides of the purported equality $\phi(\theta\delta) = \phi(\theta)\alpha(\delta) + \theta\phi(\delta)$ are bilinear functions of θ and δ. $\square$

From derivatives to homomorphisms. We now prove the following lemma.

LEMMA 9.20. *If $\phi \in Der_n$, then the map $\alpha \circ \phi : F_n \to \mathbb{Z}$ is a homomorphism.*

PROOF. For $x, y \in F_n$ we have

$$\alpha(\phi(xy)) = \alpha(\phi(x) + x\phi(y)) = \alpha(\phi(x)) + \alpha(x)\alpha(\phi(y)) = \alpha(\phi(x)) + \alpha(\phi(y)). \quad \square$$

The following lemma shows that $\alpha \circ \frac{\partial}{\partial x_a}$ is the homomorphism that counts the signed number of occurances of x_a in a word.

LEMMA 9.21. *For $1 \leqslant a \leqslant n$, the homomorphism $\alpha \circ \frac{\partial}{\partial x_a} : F_n \to \mathbb{Z}$ is the homomorphism taking x_a to 1 and x_i to 0 for $1 \leqslant i \leqslant n$ with $i \neq a$.*

PROOF. Immediate. $\square$

COROLLARY 9.22. *Consider some $w \in F_n$. Then $\alpha(\phi(w)) = 0$ for all $\phi \in Der_n$ if and only if $w \in [F_n, F_n]$.*

PROOF. Lemma 9.17 implies that $\alpha(\phi(w)) = 0$ for all $\phi \in \mathrm{Der}_n$ if and only if $\alpha(\frac{\partial}{\partial x_a}(w)) = 0$ for all $1 \leqslant a \leqslant n$. Lemma 9.21 implies that this holds if and only if $w \in [F_n, F_n]$. $\square$

Higher derivatives. An *order k free derivative* is a function $\phi : F_n \to \mathbb{Z}[F_n]$ of the form $\phi = \phi_1 \circ \phi_2 \circ \cdots \circ \phi_k$, where $\phi_i \in \mathrm{Der}_n$. Here we are extending the ϕ_i to functions $\phi_i : \mathbb{Z}[F_n] \to \mathbb{Z}[F_n]$ by linearity as above. Let Der_n^k be the set of all order k free derivatives. Define

$$\Gamma_n^k = \{w \in F_n \mid \alpha(\phi(w)) = 0 \text{ for all } \phi \in \mathrm{Der}_n^\ell \text{ with } 1 \leqslant \ell \leqslant k\}.$$

We then have the following.

LEMMA 9.23. *Consider $x \in F_n$ and $w \in \Gamma_n^k$ and $\phi \in Der_n^\ell$ for some $k, \ell \geqslant 1$ satisfying $\ell \leqslant k + 1$. Then*

$$\phi(xw) = \phi(x) + x\phi(w).$$

PROOF. The proof is by induction on ℓ. The base case $\ell = 1$ is the definition of a derivation. Now assume that $1 < \ell \leqslant k + 1$. We can write $\phi = \phi' \circ \phi''$ with $\phi' \in \mathrm{Der}_n$ and $\phi'' \in \mathrm{Der}_n^{\ell-1}$. Using our inductive hypothesis and Lemma 9.19, we have

$$\begin{aligned}
\phi(xw) &= \phi'(\phi''(xw)) \\
&= \phi'(\phi''(x) + x\phi''(w)) \\
&= \phi'(\phi''(x)) + \phi'(x)\alpha(\phi''(w)) + x\phi'(\phi''(w)) \\
&= \phi(x) + x\phi(w).
\end{aligned}$$

Here the fourth equality uses the fact that $\alpha(\phi''(w)) = 0$, which follows from the fact that $w \in \Gamma_n^k$. $\square$

Corollary 9.22 says that $\Gamma_n^1 = [F_n, F_n]$. This is in particular a normal subgroup of F_n. The following lemma generalizes this.

Lemma 9.24. *The set Γ_n^k is a normal subgroup of F_n for all $k \geqslant 1$.*

Proof. We first prove that it is a subgroup. Consider $w_1, w_2 \in \Gamma_n^k$. Our goal is to show that $w_1 w_2 \in \Gamma_n^k$. In other words, we want to show that $\alpha(\phi(w_1 w_2)) = 0$ for all $\phi \in \mathrm{Der}_n^\ell$ with $1 \leqslant \ell \leqslant k$. This follows immediately from Lemma 9.23, which implies that

$$\alpha(\phi(w_1 w_2)) = \alpha(\phi(w_1)) + \alpha(w_1)\alpha(\phi(w_2)) = 0 + 0.$$

We now prove that Γ_n^k is a normal subgroup. Consider $w \in \Gamma_n^k$ and $x \in F_n$. Our goal is to show that $xwx^{-1} \in \Gamma_n^k$. In other words, we want to show that $\alpha(\phi(xwx^{-1})) = 0$ for all $\phi \in \mathrm{Der}_n^\ell$ with $1 \leqslant \ell \leqslant k$. Define $\phi' : F_n \to \mathbb{Z}[F_n]$ via the formula

$$\phi'(z) = x^{-1}\phi(xzx^{-1})x.$$

We claim that $\phi' \in \mathrm{Der}_n^\ell$. Indeed, write $\phi = \phi_1 \circ \cdots \circ \phi_\ell$ with $\phi_i \in \mathrm{Der}_n$ for $1 \leqslant i \leqslant \ell$. Define $\phi_i' : \mathbb{Z}[F_n] \to \mathbb{Z}[F_n]$ via the formula

$$\phi_i'(z) = x^{-1}\phi_i(xzx^{-1})x.$$

For $z_1, z_2 \in F_n$ we have

$$\begin{aligned}\phi_i'(z_1 z_2) &= x^{-1}\phi_i(xz_1z_2x^{-1})x \\ &= x^{-1}(\phi_i(xz_1x^{-1}) + xz_1x^{-1}\phi_i(xz_2x^{-1}))x \\ &= \phi_i'(z_1) + z_1\phi_i'(z_2),\end{aligned}$$

so $\phi_i' \in \mathrm{Der}_n$. Since $\phi' = \phi_1' \circ \cdots \circ \phi_\ell'$, we deduce that $\phi' \in \mathrm{Der}_n^\ell$, as claimed. We now deduce that

$$\alpha(\phi(xwx^{-1})) = \alpha(x^{-1}\phi(xwx^{-1})x) = \alpha(\phi'(w)) = 0,$$

as desired. □

Remark 9.25. In fact, using the same paper of Witt [**Wit37**] that we cited when discussing the analogous fact for the Magnus expansions, Fox [**Fox53**] proved that $\Gamma_n^k = \gamma_{k+1}(F_n)$. See Lemma 9.28 below for the easy half of this.

From derivatives to homomorphisms II. We now prove the following.

Lemma 9.26. *If $\phi \in \mathrm{Der}_n^{k+1}$, then the map $\alpha \circ \phi : \Gamma_n^k \to \mathbb{Z}$ is a homomorphism.*

Proof. Consider $w_1, w_2 \in \Gamma_n^k$. Lemma 9.23 implies that

$$\alpha \circ \phi(w_1 w_2) = \alpha(\phi(w_1) + w_1\phi(w_2)) = \alpha \circ \phi(w_1) + \alpha \circ \phi(w_2). \qquad \square$$

The homomorphism constructed in Lemma 9.26 has the following property.

Lemma 9.27. *If $\phi \in \mathrm{Der}_n^{k+1}$ and $w \in \Gamma_n^k$ and $x \in F_n$, then $\alpha \circ \phi(xwx^{-1}) = \alpha \circ \phi(w)$.*

Proof. For $1 \leqslant \ell \leqslant k+1$, let S_ℓ be the subset of $\mathbb{Z}[F_n]$ consisting of all $\mathbb{Z}$-linear combinations of elements of the set

$$\{\eta(u) \mid \eta \in \mathrm{Der}_n^{\ell'} \text{ for some } 1 \leqslant \ell' < \ell \text{ and } u \in \Gamma_n^k\}.$$

By definition, we have $\alpha(s) = 0$ for all $s \in S_\ell$. We will prove that for all $1 \leqslant \ell \leqslant k+1$ and all $\eta \in \mathrm{Der}_n^\ell$, we have

$$\eta(xwx^{-1}) = \eta(x) + x\eta(w) - xwx^{-1}\eta(x) + s \tag{20}$$

for some $s \in S_\ell$. Applying this to $\eta = \phi$ and composing the result with α, it will follow that

$$\begin{aligned}\alpha(\phi(xwx^{-1})) &= \alpha(\phi(x)) + \alpha(x\phi(w)) - \alpha(xwx^{-1}\phi(x)) + \alpha(s)\\ &= \alpha(\phi(x)) + \alpha(\phi(w)) - \alpha(\phi(x)) + 0\\ &= \alpha(\phi(w)),\end{aligned}$$

as desired.

It remains to prove (20). The proof will be by induction on ℓ. The base case is $\ell = 1$. In this case, we can apply Lemma 9.18 and get that

$$\eta(xwx^{-1}) = \eta(x) + x\eta(w) + xw\eta(x^{-1}) = \eta(x) + x\eta(w) - xwx^{-1}\eta(x),$$

as claimed. Now assume that $1 < \ell \leqslant k+1$ and that the result is true for all smaller ℓ. Write $\eta = \eta' \circ \eta''$ with $\eta' \in \mathrm{Der}_n$ and $\eta'' \in \mathrm{Der}_n^{\ell-1}$. Our inductive hypothesis says that

$$\eta''(xwx^{-1}) = \eta''(x) + x\eta''(w) - xwx^{-1}\eta''(x) + s$$

for some $s \in S_{\ell-1}$. Using the fact that $\alpha(\eta''(w)) = 0$, we now apply Lemma 9.19 to see that

$$\begin{aligned}\eta(xwx^{-1}) &= \eta(x) + \eta'(x)\alpha(\eta''(w)) + x\eta(x)\\ &\qquad - \eta'(xwx^{-1})\alpha(\eta''(x)) - xwx^{-1}\eta(x) + \eta'(s)\\ &= \eta(x) + x\eta(x) - xwx^{-1}\eta(x) + s'\end{aligned}$$

with

$$s' = -\eta'(xwx^{-1})\alpha(\eta''(x)) + \eta'(s) \in S_\ell,$$

as claimed. □

Lower central series. As we said above, Fox [**Fox53**] proved that $\Gamma_n^k = \gamma_{k+1}(F_n)$. One direction of this is easy.

Lemma 9.28. *For all $k \geqslant 1$, we have $\gamma_{k+1}(F_n) \subset \Gamma_n^k$.*

Proof. The proof is by induction on k. The base case $k = 1$ is Corollary 9.22. Now assume that $k > 1$ and that the lemma is true for all smaller k. Consider $w \in \gamma_k(F_n)$ and $x \in F_n$. We must show that $[x, w] \in \Gamma_n^k$, i.e. that for all $\phi \in \mathrm{Der}_n^\ell$ with $1 \leqslant \ell \leqslant k$ we have $\alpha(\phi([x, w])) = 0$. Since $[x, w] \in \gamma_k(F_n) \subset \Gamma_n^{k-1}$, this holds if $\ell < k$. If $\ell = k$, then we can apply Lemmas 9.26 and 9.27 to see that

$$\alpha(\phi([x, w])) = \alpha(\phi(xwx^{-1})) + \alpha(\phi(w^{-1})) = \alpha(\phi(w)) - \alpha(\phi(w)) = 0,$$

as desired. □

A calculation. We now record the following calculation. Recall our convention that $[x_i, x_j] = x_i x_j x_i^{-1} x_j^{-1}$.

LEMMA 9.29. *For* $1 \leqslant a < b \leqslant n$ *and* $1 \leqslant i < j \leqslant n$, *we have*

$$\alpha(\frac{\partial}{\partial x_a} \circ \frac{\partial}{\partial x_b}([x_i, x_j])) = \begin{cases} 1 & \text{if } a = i \text{ and } b = j, \\ 0 & \text{otherwise.} \end{cases}$$

PROOF. There are several cases; we will do the case where $a = i$ and $b = j$ and leave the others to the reader. Observe that

$$\frac{\partial}{\partial x_b}([x_a, x_b]) = x_a - [x_a, x_b],$$

and hence

$$\frac{\partial}{\partial x_a}(\frac{\partial}{\partial x_b}([x_a, x_b])) = 1 - (1 - x_a x_b x_a^{-1}).$$

Applying α, we get that

$$\alpha(\frac{\partial}{\partial x_a}(\frac{\partial}{\partial x_b}([x_a, x_b]))) = 1 - (1 - 1) = 1.$$

□

The quotient of the commutator subgroup. We finally prove Theorem 9.1.

PROOF OF THEOREM 9.1. Define a set map $\rho : [F_n, F_n] \to \wedge^2 \mathbb{Z}^n$ via the formula

$$\rho(w) = \sum_{1 \leqslant a < b \leqslant n} \alpha(\frac{\partial}{\partial x_a} \circ \frac{\partial}{\partial x_b}(w))[x_a] \wedge [x_b].$$

Lemma 9.26 says that the maps $[F_n, F_n] \to \mathbb{Z}$ given by

$$w \mapsto \alpha(\frac{\partial}{\partial x_a} \circ \frac{\partial}{\partial x_b}(w))$$

are homomorphisms, so ρ is a homomorphism. We must prove that $\rho([x, y]) = [x] \wedge [y]$ for $x, y \in F_n$.

Define a set map $\eta : F_n \times F_n \to \wedge^2 \mathbb{Z}^n$ via the formula $\eta(x, y) = \rho([x, y])$. Our goal is to prove that $\eta(x, y) = [x] \wedge [y]$. We begin by proving a sequence of properties of η. As notation, for $v, w \in F_n$ we write v^w for wvw^{-1}.

- For all $x, y \in F_n$, we have $\eta(x, y) = -\eta(y, x)$. This follows from the fact that $[x, y] = [y, x]^{-1}$. This relation is reflected in $\wedge^2 \mathbb{Z}^n$ as $[x] \wedge [y] = -[y] \wedge [x]$.
- For all $x, y, z \in F_n$, we have $\eta(xz, y) = \eta(x, y) + \eta(z, y)$. Using the easily-verified commutator identity

  $$[xz, y] = [z, y]^x [x, y]$$

  and Lemma 9.27, we have

  $$\eta(xz, y) = \rho([z, y]^x) + \rho([x, y]) = \rho([z, y]) + \rho([x, y]) = \eta(x, y) + \eta(z, y),$$

  as desired. This relation is reflected in $\wedge^2 \mathbb{Z}^n$ as

  $$[xz] \wedge [y] = ([x] + [z]) \wedge [y] = [x] \wedge [y] + [z] \wedge [y].$$
- For all $x, y \in F_n$, we have $\eta(x^{-1}, y) = -\eta(x, y)$. This follows from the previous bullet point and the easy identity

  $$\eta(1, y) = \rho([1, y]) = \rho(1) = 0.$$

  This relation is reflected in $\wedge^2 \mathbb{Z}^n$ as $(-[x]) \wedge [y] = -([x] \wedge [y])$.

Using the above three facts repeatedly, we reduce the desired fact to showing that $\eta(x_i, x_j) = [x_i] \wedge [x_j]$ for $1 \leqslant i < j \leqslant n$. This follows immediately from Lemma 9.29. □

Relationship between Fox calculus and Magnus expansions. The results in this section are related to the results in §9.1 via the following lemma. Let $\psi_k : F_n \to \mathcal{A}_k(\mathbb{Z}^n)$ be the degree k Magnus expansion.

LEMMA 9.30. *For all $k \geqslant 1$, we have*

$$\psi_k(w) = 1 + \sum_{\ell=1}^{k} \Big(\sum_{1 \leqslant i_1, \ldots, i_\ell \leqslant n} (\alpha(\frac{\partial}{\partial x_{i_1}} \cdots \frac{\partial}{\partial x_{i_\ell}}(w)))[x_{i_1}][x_{i_2}] \cdots [x_{i_\ell}]\Big).$$

See [**Fox53**, §3] for the details of the proof of this.

9.3. Via group cohomology

We now sketch a final proof of Theorem 9.1. The main tool is the following theorem concerning group homology. Recall that if G is a group, then $\mathrm{H}_k(G;\mathbb{Z})$ is defined to be the k^{th} homology group of a $K(G,1)$. Also, if M is an abelian group on which G acts, then the *coinvariants* of M, denoted M_G, is the quotient of M by the submodule $\langle m - g(m) \mid m \in M, \, g \in G\rangle$.

THEOREM 9.31. *If*

$$1 \longrightarrow K \longrightarrow G \longrightarrow Q \longrightarrow 1$$

is a short exact sequence of groups, then there exists a 5-term exact sequence

$$\mathrm{H}_2(G;\mathbb{Z}) \longrightarrow \mathrm{H}_2(Q;\mathbb{Z}) \longrightarrow (\mathrm{H}_1(K;\mathbb{Z}))_G \longrightarrow \mathrm{H}_1(G;\mathbb{Z}) \longrightarrow \mathrm{H}_1(Q;\mathbb{Z}) \longrightarrow 0.$$

In the statement of Theorem 9.31, the action of G on $\mathrm{H}_1(K;\mathbb{Z})$ is induced by the conjugation action of G on K. Theorem 9.31 (or, rather, a dual statement in cohomology) was first proven by Hochschild–Serre [**HS53**, Theorem 2] as an application of the Hochschild–Serre spectral sequence in group homology; see [**Bro94**, Chapter VII.6] for a textbook exposition of this proof and [**Bro94**, Exercise II.5.6] and [**Coc85**] for alternate proofs.

PROOF OF THEOREM 9.1 (SKETCH). Consider the short exact sequence

$$1 \longrightarrow [F_n, F_n] \longrightarrow F_n \longrightarrow \mathbb{Z}^n \longrightarrow 1.$$

The associated 5-term exact sequence given by Theorem 9.31 is of the form

$$\mathrm{H}_2(F_n;\mathbb{Z}) \longrightarrow \mathrm{H}_2(\mathbb{Z}^n;\mathbb{Z}) \longrightarrow (\mathrm{H}_1([F_n,F_n]))_{F_n} \longrightarrow \mathrm{H}_1(F_n;\mathbb{Z}) \longrightarrow \mathrm{H}_1(\mathbb{Z}^n;\mathbb{Z}) \longrightarrow 0.$$

Since F_n is a free group, we have $\mathrm{H}_2(F_n;\mathbb{Z}) = 0$. Also, the map $\mathrm{H}_1(F_n;\mathbb{Z}) \to \mathrm{H}_1(\mathbb{Z}^n;\mathbb{Z})$ is an isomorphism. Finally, using the fact that the n-torus $\mathbb{T}^n$ is a $K(\mathbb{Z}^n,1)$ we have that $\mathrm{H}_2(\mathbb{Z}^n;\mathbb{Z}) \cong \wedge^2\mathbb{Z}^n$. We conclude that

$$\wedge^2\mathbb{Z}^n \cong (\mathrm{H}_1([F_n,F_n]))_{F_n}.$$

The map $\rho : [F_n,F_n] \to \wedge^2\mathbb{Z}^n$ is then the composition

$$[F_n,F_n] \to [F_n,F_n]^{\text{ab}} = \mathrm{H}_1([F_n,F_n];\mathbb{Z}) \to (\mathrm{H}_1([F_n,F_n];\mathbb{Z}))_{F_n} \cong \wedge^2\mathbb{Z}^n.$$

The claimed description of ρ follows from a careful examination of the maps involved in the proof of Theorem 9.31. □

Bibliography

[Bae27] R. Baer, *Kurventypen auf Flächen.*, J. Reine Angew. Math. **156** (1927), 231–246 (German). (cited on page 5.)

[Bae28] ———, *Isotopie von Kurven auf orientierbaren, geschlossenen Flächen und ihr Zusamenhang mit der topologischen Deformation der Flächen.*, J. Reine Angew. Math. **159** (1928), 101–116 (German). (cited on page 5.)

[BBM10] Mladen Bestvina, Kai-Uwe Bux, and Dan Margalit, *The dimension of the Torelli group*, J. Amer. Math. Soc. **23** (2010), no. 1, 61–105. MR 2552249 (2011b:20109) (cited on pages 43, 52, 65, and 73.)

[Bir69] Joan S. Birman, *Mapping class groups and their relationship to braid groups*, Comm. Pure Appl. Math. **22** (1969), 213–238. MR 0243519 (39 #4840) (cited on page 12.)

[Bir71] ———, *On Siegel's modular group*, Math. Ann. **191** (1971), 59–68. MR 0280606 (43 #6325) (cited on pages 29 and 41.)

[Bre97] Glen E. Bredon, *Topology and geometry*, Graduate Texts in Mathematics, vol. 139, Springer-Verlag, New York, 1997, Corrected third printing of the 1993 original. MR 1700700 (2000b:55001) (cited on pages 7 and 48.)

[Bro94] Kenneth S. Brown, *Cohomology of groups*, Graduate Texts in Mathematics, vol. 87, Springer-Verlag, New York, 1994, Corrected reprint of the 1982 original. MR 1324339 (96a:20072) (cited on page 97.)

[Bur90] H. Burkhardt, *Grundzüge einer allgemeinen Systematik der hyperelliptischen Functionen I. Ordnung. Nach Vorlesungen von F. Klein.*, Math. Ann. **35** (1890), 198–296 (German). (cited on page 19.)

[CFL58] K.-T. Chen, R. H. Fox, and R. C. Lyndon, *Free differential calculus. IV. The quotient groups of the lower central series*, Ann. of Math. (2) **68** (1958), 81–95. MR 0102539 (21 #1330) (cited on page 90.)

[CG66] Alfred Clebsch and Paul Gordan, *Theorie der abelschen funktionen*, B. G. Teubner, Leipzig, 1966 (German). (cited on page 19.)

[Coc85] Tim D. Cochran, *A topological proof of Stallings' theorem on lower central series of groups*, Math. Proc. Cambridge Philos. Soc. **97** (1985), no. 3, 465–472. MR 778680 (86e:57002) (cited on page 97.)

[Deh38] M. Dehn, *Die Gruppe der Abbildungsklassen. (Das arithmetische Feld auf Flächen.).*, Acta Math. **69** (1938), 135–206 (German). (cited on pages 9, 10, and 15.)

[Eps66] D. B. A. Epstein, *Curves on 2-manifolds and isotopies*, Acta Math. **115** (1966), 83–107. MR 0214087 (35 #4938) (cited on page 5.)

[Fen50] W. Fenchel, *Remarks on finite groups of mapping classes*, Mat. Tidsskr. B. **1950** (1950), 90–95. MR 0038069 (12,349f) (cited on page 30.)

[FM12] Benson Farb and Dan Margalit, *A primer on mapping class groups*, Princeton Mathematical Series, vol. 49, Princeton University Press, Princeton, NJ, 2012. MR 2850125 (2012h:57032) (cited on pages 5, 6, 9, 10, 12, 13, 14, 15, 30, 34, 48, and 67.)

[Fox53] Ralph H. Fox, *Free differential calculus. I. Derivation in the free group ring*, Ann. of Math. (2) **57** (1953), 547–560. MR 0053938 (14,843d) (cited on pages 91, 94, 95, and 97.)

[Hir94] Morris W. Hirsch, *Differential topology*, Graduate Texts in Mathematics, vol. 33, Springer-Verlag, New York, 1994, Corrected reprint of the 1976 original. MR 1336822 (96c:57001) (cited on pages 6 and 9.)

[HM12] Allen Hatcher and Dan Margalit, *Generating the Torelli group*, Enseign. Math. (2) **58** (2012), no. 1-2, 165–188. MR 2985015 (cited on pages 29 and 43.)

[HS53] G. Hochschild and J.-P. Serre, *Cohomology of group extensions*, Trans. Amer. Math. Soc. **74** (1953), 110–134. MR 0052438 (14,619b) (cited on page 97.)

[HT80] A. Hatcher and W. Thurston, *A presentation for the mapping class group of a closed orientable surface*, Topology **19** (1980), no. 3, 221–237. MR 579573 (81k:57008) (cited on page 15.)

[Hum79] Stephen P. Humphries, *Generators for the mapping class group*, Topology of low-dimensional manifolds (Proc. Second Sussex Conf., Chelwood Gate, 1977), Lecture Notes in Math., vol. 722, Springer, Berlin, 1979, pp. 44–47. MR 547453 (80i:57010) (cited on page 10.)

[Hur93] A. Hurwitz, *Ueber algebraische Gebilde mit eindeutigen Transformationen in sich.*, Math. Ann. **41** (1893), 403–442 (German). (cited on page 31.)

[Iva02] Nikolai V. Ivanov, *Mapping class groups*, Handbook of geometric topology, North-Holland, Amsterdam, 2002, pp. 523–633. MR 1886678 (2003h:57022) (cited on pages 5, 9, and 10.)

[Joh79] Dennis Johnson, *Homeomorphisms of a surface which act trivially on homology*, Proc. Amer. Math. Soc. **75** (1979), no. 1, 119–125. MR 529227 (80h:57008) (cited on pages 15, 29, and 41.)

[Joh80] ———, *Conjugacy relations in subgroups of the mapping class group and a group-theoretic description of the Rochlin invariant*, Math. Ann. **249** (1980), no. 3, 243–263. MR 579104 (82a:57007) (cited on pages 32, 34, and 36.)

[Joh83] ———, *The structure of the Torelli group. I. A finite set of generators for $\mathcal{I}$*, Ann. of Math. (2) **118** (1983), no. 3, 423–442. MR 727699 (85a:57005) (cited on pages 12, 30, and 41.)

[Joh85a] ———, *The structure of the Torelli group. II. A characterization of the group generated by twists on bounding curves*, Topology **24** (1985), no. 2, 113–126. MR 793178 (86i:57011) (cited on page 41.)

[Joh85b] ———, *The structure of the Torelli group. III. The abelianization of $\mathcal{T}$*, Topology **24** (1985), no. 2, 127–144. MR 793179 (87a:57016) (cited on page 41.)

[Ker83] Steven P. Kerckhoff, *The Nielsen realization problem*, Ann. of Math. (2) **117** (1983), no. 2, 235–265. MR 690845 (85e:32029) (cited on page 30.)

[Lef26] Solomon Lefschetz, *Intersections and transformations of complexes and manifolds*, Trans. Amer. Math. Soc. **28** (1926), no. 1, 1–49. MR 1501331 (cited on page 31.)

[Lic64] W. B. R. Lickorish, *A finite set of generators for the homeotopy group of a 2-manifold*, Proc. Cambridge Philos. Soc. **60** (1964), 769–778. MR 0171269 (30 #1500) (cited on pages 9 and 10.)

[Mag35] Wilhelm Magnus, *Beziehungen zwischen Gruppen und Idealen in einem speziellen Ring*, Math. Ann. **111** (1935), no. 1, 259–280. MR 1512992 (cited on pages 87 and 90.)

[McC75] James McCool, *Some finitely presented subgroups of the automorphism group of a free group*, J. Algebra **35** (1975), 205–213. MR 0396764 (53 #624) (cited on page 15.)

[Mes92] Geoffrey Mess, *The Torelli groups for genus 2 and 3 surfaces*, Topology **31** (1992), no. 4, 775–790. MR 1191379 (93k:57003) (cited on pages 28, 30, 41, 43, 65, and 73.)

[Mey76] Mark D. Meyerson, *Representing homology classes of closed orientable surfaces*, Proc. Amer. Math. Soc. **61** (1976), no. 1, 181–182 (1977). MR 0425967 (54 #13916) (cited on pages 19 and 20.)

[MKS76] Wilhelm Magnus, Abraham Karrass, and Donald Solitar, *Combinatorial group theory*, revised ed., Dover Publications, Inc., New York, 1976, Presentations of groups in terms of generators and relations. MR 0422434 (54 #10423) (cited on pages 87 and 90.)

[MM86] Darryl McCullough and Andy Miller, *The genus 2 Torelli group is not finitely generated*, Topology Appl. **22** (1986), no. 1, 43–49. MR 831180 (87h:57015) (cited on page 30.)

[MM09] Dan Margalit and Jon McCammond, *Geometric presentations for the pure braid group*, J. Knot Theory Ramifications **18** (2009), no. 1, 1–20. MR 2490001 (2010a:20086) (cited on page 15.)

[MP78] William H. Meeks, III and Julie Patrusky, *Representing homology classes by embedded circles on a compact surface*, Illinois J. Math. **22** (1978), no. 2, 262–269. MR 0474304 (57 #13951) (cited on page 19.)

[MP10] Justin Malestein and Andrew Putman, *On the self-intersections of curves deep in the lower central series of a surface group*, Geom. Dedicata **149** (2010), 73–84. MR 2737679 (2012f:57003) (cited on page 90.)

[Nie43] Jakob Nielsen, *Abbildungsklassen endlicher Ordnung*, Acta Math. **75** (1943), 23–115. MR 0013306 (7,137a) (cited on page 30.)

[Pow78] Jerome Powell, *Two theorems on the mapping class group of a surface*, Proc. Amer. Math. Soc. **68** (1978), no. 3, 347–350. MR 0494115 (58 #13045) (cited on pages 29 and 41.)

[Put07] Andrew Putman, *Cutting and pasting in the Torelli group*, Geom. Topol. **11** (2007), 829–865. MR 2302503 (2008c:57049) (cited on pages 27, 29, 32, 57, and 62.)

[Put09] ———, *An infinite presentation of the Torelli group*, Geom. Funct. Anal. **19** (2009), no. 2, 591–643. MR 2545251 (2010k:57069) (cited on page 15.)

[Put12] ———, *Small generating sets for the Torelli group*, Geom. Topol. **16** (2012), no. 1, 111–125. MR 2872579 (2012k:57003) (cited on page 30.)

[Sch76] James A. Schafer, *Representing homology classes on surfaces*, Canad. Math. Bull. **19** (1976), no. 3, 373–374. MR 0454995 (56 #13236) (cited on page 19.)

[Ser80] Jean-Pierre Serre, *Trees*, Springer-Verlag, Berlin, 1980, Translated from the French by John Stillwell. MR 607504 (82c:20083) (cited on pages 70 and 71.)

[Ser92] ———, *Lie algebras and Lie groups*, second ed., Lecture Notes in Mathematics, vol. 1500, Springer-Verlag, Berlin, 1992, 1964 lectures given at Harvard University. MR 1176100 (93h:17001) (cited on page 90.)

[SW79] Peter Scott and Terry Wall, *Topological methods in group theory*, Homological group theory (Proc. Sympos., Durham, 1977), London Math. Soc. Lecture Note Ser., vol. 36, Cambridge Univ. Press, Cambridge, 1979, pp. 137–203. MR 564422 (81m:57002) (cited on pages 70 and 71.)

[vdB03] Barbara van den Berg, *On the abelianization of the torelli group*, University of Utrecht, 2003. (cited on pages 57 and 62.)

[Wit37] Ernst Witt, *Treue Darstellung Liescher Ringe.*, J. Reine Angew. Math. **177** (1937), 152–160 (German). (cited on pages 90 and 94.)