Jean-Luc Thiffeault

Lectures on Braids and Dynamics

From taffy pullers to data analysis

Instituto de Matemática
Universidade Federal do Rio de Janeiro, May 3–12, 2022

Contents

Lecture 1
Topological dynamics on the torus

Mmm... donuts.
— Homer Simpson

In this lecture we use the torus to illustrate the basic ideas behind the topological classification of mappings. We introduce the mapping class group of the torus and investigate its properties.

1.1 Diffeomorphisms of the torus

One of the simplest surfaces to study is the *torus*, often denoted T^2. Its 3D representation is depicted in Fig. 1.1a, but for us it will often be more convenient to use the 'flattened' view of Fig. 1.1b. The arrows in Fig. 1.1b indicate the edges that are identified with each other.

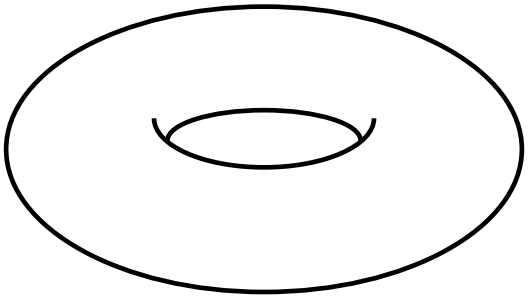

Fig. 1.1a The humble torus.

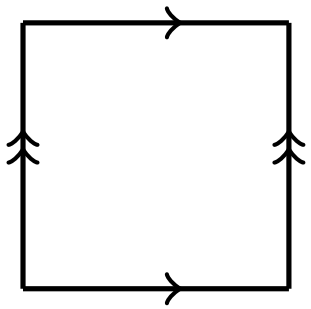

Fig. 1.1b A different view of the torus.

In this lecture, the central question we want to answer is: what do maps of the torus to itself look like? The relevant maps are *homeomorphisms*: continuous, invertible maps whose inverse is also continous. In fact it costs us nothing to use instead *diffeomorphisms*:

smooth, invertible maps whose inverse is also smooth. (See Farb and Margalit [11] for the reason why homeomorphisms can always be promoted to diffeomorphisms in the context of topology of surfaces.) The set of such maps forms a group, $\mathrm{Diff}(T^2)$, where the group operation is simply composition of maps. In fact, we shall focus solely on $\mathrm{Diff}^+(T^2)$, the group of orientation-preserving diffeomorphisms.

We will classify diffeomorphisms a surface S up to *isotopy*, an equivalence relation denoted by $\simeq$. Two diffeomorphisms are isotopic if they can be continuously 'deformed' into each other. More precisely, two diffeomorphisms ϕ and $\psi \in \mathrm{Diff}^+(S)$ are *isotopic* if there exists a homotopy

$$H : S \times I \to S \quad \text{with} \quad H(S \times \{0\}) = \phi, \quad H(S \times \{1\}) = \psi, \tag{1.1}$$

such that $H(S \times \{t\}) \in \mathrm{Diff}^+(S)$ for all $t \in I = [0,1]$. Note that

$$\phi \simeq \psi \qquad \Longleftrightarrow \qquad \phi = \chi \circ \psi \quad \text{with} \quad \chi \simeq \mathrm{id} \tag{1.2}$$

that is, χ is isotopic to the identity map. We write $\mathrm{Diff}_0(S)$ for the subgroup of maps isotopic to the identity.

We define the *mapping class group* of a surface S by

$$\mathrm{MCG}(S) := \mathrm{Diff}^+(S)/\mathrm{Diff}_0(S). \tag{1.3}$$

This is the central object that we wish to study: the mapping class group describes self-maps of a surface, up to isotopy. In that sense it distills the essential topological information about such maps.

1.2 The mapping class group of the torus

Let us start with a bit of terminology. A *curve* on a surface S is the image of a continuous map $f : I \to S$, where $I = [0,1]$. An *oriented* curve carries an orientation as well. For a *simple* curve f is injective in $(0,1)$, so the curve does not intersect itself. A *closed* curve has $f(0) = f(1)$. An *essential* curve is not contractible to a point or a boundary component. We shall often use *loop* for an oriented closed curve.

Figure 1.2 shows the action of a simple diffeomorphism on two loops that intersect at one point. Note that the number of intersections cannot change, since the map is one-to-one. This picture illustrates the induced action of the map ϕ on the *first homotopy group* of the torus, $\pi_1(T^2, x_0)$:

$$\phi : T^2 \to T^2 \qquad \text{induces} \qquad \phi_* : \pi_1(T^2, x_0) \to \pi_1(T^2, x_0). \tag{1.4}$$

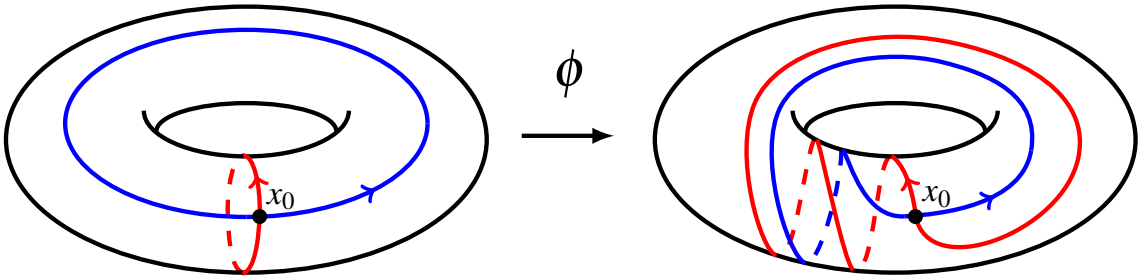

Fig. 1.2 The action on loops of a mapping of the torus.

The first homotopy group is also called the *fundamental group* of T^2 with *base point* x_0. The fundamental group contains equivalence classes of loops on a surface, where the equivalence is homotopy of paths (continuous deformations fixing x_0). The closed curves begin and end at x_0, the intersection of the curves in Fig. 1.2. We assume for simplicity that ϕ preserves the base point ($\phi(x_0) = x_0$). The fundamental group of the torus is $\mathbb{Z}^2$ [18], so it is actually isomorphic to $H_1(T^2,\mathbb{Z})$, the first homotopy group with coefficients in $\mathbb{Z}$. Hence, we can represent the action of ϕ_* by a matrix in $\mathrm{SL}_2(\mathbb{Z})$. The matrix must have determinant one, since ϕ^{-1} must also induce an action ϕ_*^{-1} on $\mathbb{Z}^2$. (A negative determinant is ruled out, since ϕ preserves orientation.)

The striking result is that the map $\sigma : \mathrm{MCG}(T^2) \to \mathrm{SL}_2(\mathbb{Z})$ given by action on first homology is an isomorphism. To put it another way, the mapping class group of the torus is given entirely by the action of diffeomorphisms on $H_1(T^2,\mathbb{Z})$. First we show that σ is well-defined: given a mapping $\phi \in \mathrm{Diff}^+(T^2)$, any other mapping isotopic to ϕ will act on $H_1(T^2,\mathbb{Z})$ with a matrix $\chi_* \circ \phi_*$, where $\chi_* \in \mathrm{Diff}_0(T^2)$. But $\chi_* = \mathrm{id}$, since the isotopy used to take χ back to the identity map will also take loops back to their initial path, leaving any loop unchanged. Hence all elements of the isotopy class of ϕ act the same way on $H_1(T^2,\mathbb{Z})$.

Next, we show that σ is surjective. Given any element M of $\mathrm{SL}_2(\mathbb{Z})$, we can construct an orientation-preserving linear diffeomorphism of $\mathbb{R}^2$. Taking the torus as the unit square with edges identified, we can see that this linear map induces a diffeomorphism of the torus. (This requires showing equivariance with respect to the deck transformation of the covering space.) This diffeomorphism induces an action on $H_1(T^2,\mathbb{Z})$ with matrix M.

The injectivity of σ is the harder part: to show that two non-isotopic maps must act differently on $H_1(T^2,\mathbb{Z})$. Since this is more technical, we refer the reader to Farb and Margalit [11, p. 53].

1.3 Classification of $\mathrm{MCG}(T^2)$

In Section 1.2, we've shown that an element of $\mathrm{MCG}(T^2)$ can be written as a matrix in $\mathrm{SL}_2(\mathbb{Z})$,

$$M = \begin{pmatrix} a & b \\ c & d \end{pmatrix}, \qquad ad - bc = 1. \tag{1.5}$$

The characteristic polynomial of M is

$$p(x) = x^2 - \tau x + 1, \qquad \tau := a + d, \tag{1.6}$$

where τ is the trace of M. We will now classify the possible types of M by examining how powers of M behave. We will make use of the characteristic polynomial $p(x)$.

The eigenvalues of M are $x_\pm = \frac{1}{2}(\tau \pm \sqrt{\tau^2 - 4})$ (with $x_+ x_- = 1$), suggesting that $|\tau| = 2$ plays a special role. In fact low values of $|\tau|$ are important, and we will increase its value gradually in our analysis.

1.3.1 The elliptic case

Recall the Cayley–Hamilton theorem, which says that a matrix is a root of its characteristic polynomial: $p(M) = 0$. (Constants are interpreted as multiples of the identity matrix, I.)

Consider first $|\tau| < 2$. Since τ is an integer, this means $\tau = -1, 0, 1$. If $\tau = 0$, then by the Cayley–Hamilton theorem $M^2 + I = 0$, so that $M^4 = I$. If $\tau = \pm 1$, then by the same theorem $M^2 \mp M + I = 0$, so $M^2 = \pm M - I$. Multiplying both sides by M, we obtain $M^3 = \pm M^2 - M = \pm(\pm M - I) - M = \mp I$. Hence, $M^6 = I$ for $|\tau| = 1$. The least common multiple of 4 and 6 is 12, so we conclude that $M^{12} = I$ for $|\tau| < 2$.

If $|\tau| = 2$ and $b = c = 0$, we have $M = \pm I$, so that $M^2 = I$.

Putting these together, we conclude that if $|\tau| < 2$ or $M = \pm I$, then $M^{12} = I$. This behavior is called *finite order* and the matrix M is *elliptic*: after a while (at most 12 applications), powers of M repeat themselves. Since M acts on homological generators, these must all return to their initial configuration. It follows that ϕ^{12} is isotopic to the identity. Note that this does not imply that ϕ^{12} itself is the identity: in fact ϕ could be a complicated map with chaotic orbits. But topologically, it is indistinguishable from the identity.

1.3.2 The parabolic case

Building on Section 1.3.1, let us consider $\tau = \pm 2$, but assume that $|M| \neq I$. The eigenvalues of M are degenerate and both equal to $\tau/2 = \pm 1$. The Cayley–Hamilton theorem says that $(M \mp I)^2 = 0$, or $M = \pm I + N$, where $N \neq 0$ is *nilpotent* ($N^2 = 0$). Such a matrix is called *parabolic*.

It is a simple exercise to show that $\vec{e} = ((a - \frac{1}{2}\tau) \;\; c)^T$ is an eigenvector (making use of $\tau/2 = 2/\tau$), unless $c = 0$ in which case we use instead $\vec{e} = (b \;\; (d - \frac{1}{2}\tau))^T$. The matrix M has only one eigenvector.

Thus, M leaves invariant a loop $\vec{e}$ and its multiples, possibly reversing their orientation if $\tau = -2$. (We can make $\vec{e}$ *primitive* by dividing by the greatest common divisor of its two entries. This means that $\vec{e}$ does not trace the same path more than once.) Every other loop is affected by M. This is in contrast to the elliptic (finite-order) case, where either every loop is untouched by M^p, or they all are. The invariant $\vec{e}$ corresponds to an equivalence class of *reducing curves*.

Two distinguished parabolic elements of MCG(T^2) play a special role:

$$T_1 = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}, \qquad T_2 = \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix}. \tag{1.7}$$

These have, respectively, a unique eigenvector $(0 \;\; 1)^T$ and $(1 \;\; 0)^T$. The first matrix, when acting on the loop $(p \;\; q)^T$, gives

$$\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} p \\ q \end{pmatrix} = \begin{pmatrix} p \\ p+q \end{pmatrix}. \tag{1.8}$$

It has changed the *twist* in the second component (for $p \neq 0$). This action is illustrated in Fig. 1.3, using the standard basis where $(1 \;\; 0)^T$ is a horizontal loop and $(0 \;\; 1)^T$ is a vertical loop. The linear diffeomorphisms constructed from T_1 and T_2 are called *Dehn twists* along the curves $(0 \;\; 1)^T$ and $(1 \;\; 0)^T$, respectively.

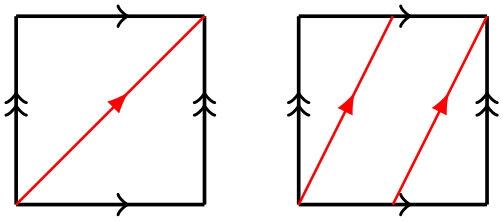

Fig. 1.3 The torus loop $(1 \;\; 1)^T$ (left) and its image $(1 \;\; 2)^T$ (right) under the action of the Dehn twist (1.8).

The two matrices (1.7) *generate* the group $\mathrm{SL}_2(\mathbb{Z})$, in the sense that any element can be written as a product of these two matrices and their inverses, possibly repeated.

1.3.3 The hyperbolic case

Finally, we come to the most interesting case: $|\tau| > 2$. In that case the two eigenvalues $x_\pm = \frac{1}{2}(\tau \pm \sqrt{\tau^2-4})$ are real. We call $\lambda = \frac{1}{2}(|\tau| + \sqrt{\tau^2-4}) > 1$ the *dilatation*[1] of the mapping class corresponding to ϕ. ($\log\lambda$ is called the *topological entropy*.) The two eigenvalues can be written $\pm\{\lambda, \lambda^{-1}\}$, where the sign is the same as τ. It is a standard result that λ is irrational for such *hyperbolic* matrices. For instance, we can easily show that the continued fraction expansion of λ (for $\tau > 2$) is $[\tau-1; 1, \tau-2, 1, \tau-2, \dots]$, with period 2.

The eigenvectors of M can be written

$$\begin{pmatrix} \pm\lambda - d \\ c \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} \pm\lambda^{-1} - d \\ c \end{pmatrix} \tag{1.9}$$

respectively with eigenvalue $\pm\lambda$ and $\pm\lambda^{-1}$. The slope of these vectors is irrational, which means that no actual loop remains invariant (or is multiplied by a constant) under the action of M. But in fact when we act repeatedly with M on some nonzero loop, the loop asympotically approaches the first eigenvector above, called the *unstable direction*. The other eigenvector, the *stable direction*, is approached under repeated iteration of M^{-1}.

If we start at any point in T^2 and draw a straight curve parallel to either eigenvector, this will give a straight line with irrational slope. The curve will thus wind around the torus but never repeat itself, and will be dense on the torus. The M-invariant object correspondig to the union of all such curves is called *the unstable foliation*, $\mathscr{F}_\mathrm{u}$, when the curves are parallel to the unstable direction (and similarly for the stable foliation, $\mathscr{F}_\mathrm{s}$). A single curve in $\mathscr{F}_\mathrm{u}$ or $\mathscr{F}_\mathrm{s}$ is called a *leaf* of the foliation. Under the action of M, a leaf of $\mathscr{F}_\mathrm{u}$ is stretched by λ, while it is contracted by λ^{-1} for $\mathscr{F}_\mathrm{s}$.

When we construct a map by taking the linear action of M on $\mathbb{R}^2$ and projecting on the torus, we obtain an *Anosov diffeomorphism*. The most famous example is *Arnold's cat map*, with matrix

$$\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}, \tag{1.10}$$

with $\lambda = \frac{1}{2}(3+\sqrt{5}) = 2.61803\dots$, the square of the Golden Ratio $\varphi = \frac{1}{2}(1+\sqrt{5}) = 1.61803\dots$. The action of this map on a loop is shown in Fig. 1.4. The number of windings around each direction of the torus increases exponentially. After many iterations, the number of windings is multiplied by λ at each iteration. The curve approaches exponentially a leaf of the unstable foliation, $\mathscr{F}_\mathrm{u}$.

[1] This is also called the dilation, stretch factor, expansion constant, or growth.

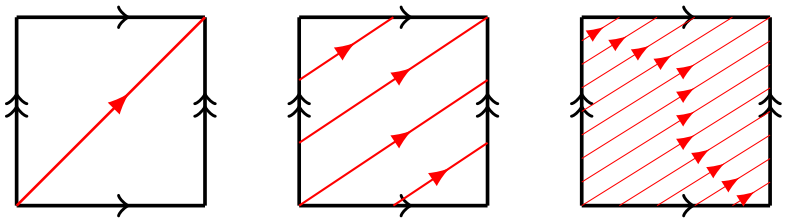

Fig. 1.4 The torus loop $(1\ 1)^T$ (left) and its images $(3\ 2)^T$ (center) and $(8\ 5)^T$ (right) under the action of the Anosov (1.10). The slope of the loop rapidly approaches the inverse of the Golden Ratio. In the last frame, the slope is $\frac{5}{8} = 0.625$, whereas $\phi^{-1} = 0.618\ldots$.

Lecture 2
The 3-rod stirring device

Pinch a torus, and you'll get two spheres.
— Anonymous

In Lecture 1 we discussed the mapping class group of the torus, and showed that it contains three types of elements: finite-order, parabolic, and Anosov. Now we will see how the abstract torus can be mapped onto the sphere with some special points. This will open up the analysis of *three-rod stirring devices*, such as the taffy puller shown in Fig. 2.1a–2.1b.

Fig. 2.1a A taffy-pulling device (from Finn and Thiffeault [12]).

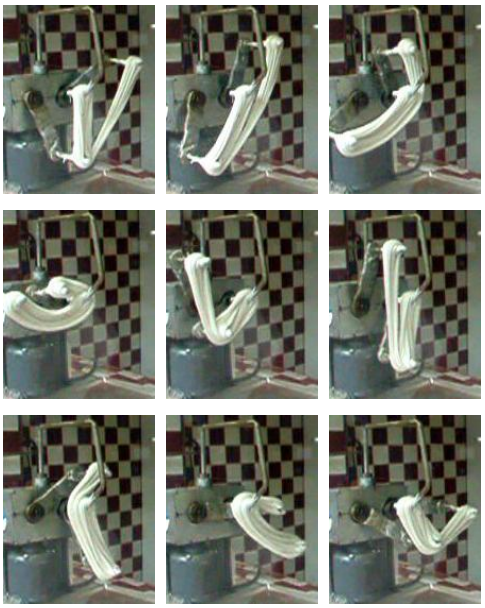

Fig. 2.1b Snapshots over a full period of operation, with taffy (from Finn and Thiffeault [12]).

2.1 From the torus to the sphere

Consider the linear map $\iota : T^2 \to T^2$, defined by $\iota(x) = -x$, where we regard the torus as the periodic unit square. This map is called the *hyperelliptic involution*, since $\iota^2 = \mathrm{id}$. It has four fixed points:

$$p_0 = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \qquad p_1 = \begin{pmatrix} \frac{1}{2} \\ 0 \end{pmatrix}, \qquad p_2 = \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix}, \qquad p_3 = \begin{pmatrix} 0 \\ \frac{1}{2} \end{pmatrix}, \tag{2.1}$$

since $-\frac{1}{2}$ is identified with $\frac{1}{2}$ because of periodicity. These are the *Weierstrass points* of ι. The map ι rotates the unit square about its center by 180°. It is in the *center* of $\mathrm{MCG}(T^2)$, since it commutes with all its elements.

Now let us construct the quotient space

$$S = T^2/\iota, \tag{2.2}$$

that is, the torus where points related by the involution are identified with each other. The map $\pi : T^2 \to T^2/\iota$ is the canonical projection. Figure 2.2a shows how the points on the torus are related to each other through ι and the identifications in Fig. 1.1b. In Fig. 2.2b we discard a redundant half of the surface.

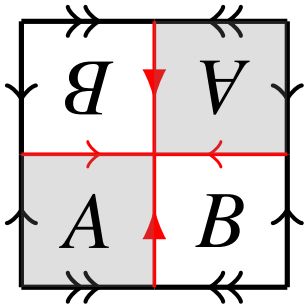

Fig. 2.2a Identifications of regions on T^2 under the map ι. The two red lines are invariant.

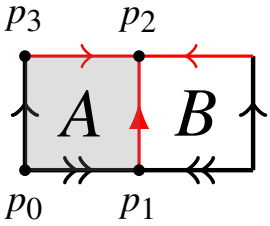

Fig. 2.2b The surface $S = T^2/\iota$, with the four Weierstrass points shown.

Next, we zip along along the identified edges between p_0 and p_1 and between p_2 and p_3, to obtain the surface in Fig. 2.3a. There is only one identification left, between the left and right boundaries of that surface. If we glue these together, we obtain a surface shaped like an American football, with p_0 and p_3 at the tips. This surface has the same topology as a sphere, with four 'distinguished' points correponsing to the Weierstrass points. This final surface is shown in Fig. 2.3b.

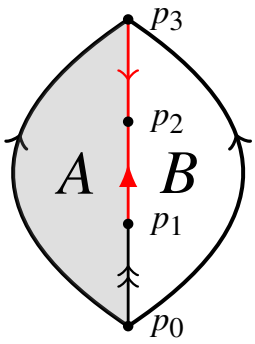

Fig. 2.3a Zipped surface.

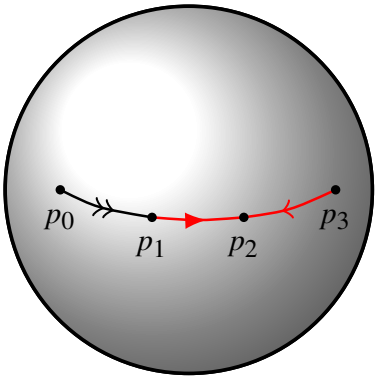

Fig. 2.3b $S = T^2/\iota$ is a sphere with four punctures, denoted $S_{0,4}$.

The four Weierstrass points can be interpreted as punctures on the sphere, so the surface in Fig. 2.3b is often denoted $S_{0,4}$: a surface of genus 0 with 4 punctures. The interpretation of the four points as punctures comes from the fact that homotopy classes of closed curves on T^2 project down under ι to homotopy classes of closed curves on $S_{0,4}$. In Fig. 2.4a we shown two close curves on T^2 that intersect once. Their image on $S_{0,4}$ is two closed curves that intersect twice (Fig. 2.4b). In fact there is a bijection between homotopy classes of

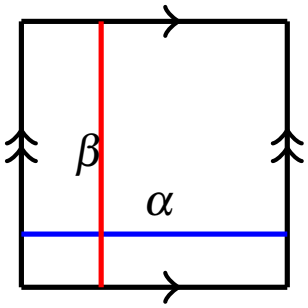

Fig. 2.4a Two intersecting closed curves on the torus.

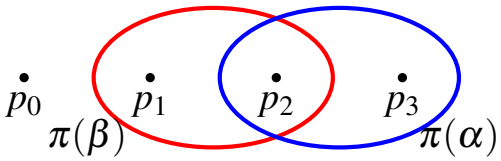

Fig. 2.4b The image of the two closed curves in Fig. 2.4a on the sphere $S_{0,4}$ (sphere not shown). The curves intersect twice. Here $\pi : T^2 \to S_{0,4}$ is the canonical projection map.

closed curves on T^2 and those on $S_{0,4}$ [11, p. 55]. The preimage of a closed curve on $S_{0,4}$ is a 'double closed curve' on T^2, since given a curve α on T^2, both α and $\iota(\alpha)$ map to the same curve on $S_{0,4}$. Note that this construction does not work for oriented curves, since the preimage of an oriented curve on $S_{0,4}$ gives two disjoint identical curves with opposite orientation on T^2 (related by ι), and therefore not a well-defined homotopy class.

2.2 The mapping class groups of $S_{0,4}$ and D_3

Mapping classes on T^2 represented by ϕ and $\iota\phi$ project down to a well-defined mapping class on $S_{0,4}$. The elements of this group are in $\mathrm{SL}_2(\mathbb{Z})/\iota =: \mathrm{PSL}_2(\mathbb{Z})$, the *projective* version of the linear group. This means that the mapping class group of $S_{0,4}$ has 'fewer' elements than that of T^2, but it gains a few more because it is possible to interchange the four punctures on the sphere pairwise through rotation, which lifts to the identity map on the (unpunctured) torus. These rotations by 180° are the two hyperelliptic involutions of $S_{0,4}$. The full mapping class group is thus [11, p. 56]

$$\mathrm{MCG}(S_{0,4}) \approx \mathrm{PSL}_2(\mathbb{Z}) \ltimes (\mathbb{Z}_2 \times \mathbb{Z}_2). \tag{2.3}$$

The operator $\ltimes$ is a *semidirect product* of the group $\mathrm{PSL}_2(\mathbb{Z})$ acting on the group $\mathbb{Z}_2 \times \mathbb{Z}_2$, the latter being the group generated by the two hyperellitic involutions of $S_{0,4}$.

We shall not dwell on this semidirect product structure, because for applications we are actually interested in a slightly different surface: D_3, the *disk* with 3 punctures. This is topologically the same as $S_{0,4}$ if we remove a small disk around one of the punctures, and then stretch this disk to make the outer boundary of D_3 (Fig. 2.5).

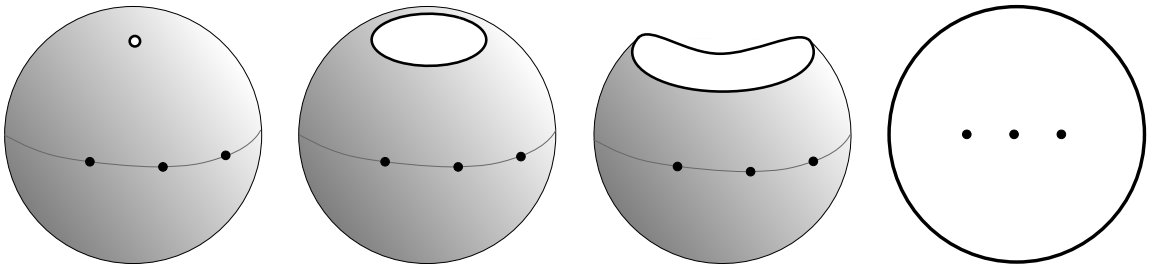

Fig. 2.5 After removing a small disk around one of the punctures (left), $S_{0,4}$ is turned into the disk D_3 by stretching the boundary (right).

A consequence of replacing a puncture by a boundary is that it is no longer equivalent to the other punctures, so diffeomorphisms cannot interchange the boundary with any puncture. This rules out the hyperelliptic involutions of $S_{0,4}$ as allowable diffeomorphisms, so we lose the factor $\mathbb{Z}_2 \times \mathbb{Z}_2$ in Eq. (2.3).

To fully characterize $\mathrm{MCG}(D_3)$, we must specify how diffeomorphisms of D_3 to itself treat the boundary, ∂D. There are two options: fix ∂D pointwise or setwise. If we fix it pointwise, a diffeomorphism can 'twist' around the boundary by rotating it by multiples of 2π. Isotopies also fix the boundary pointwise, so these cannot be untwisted. The mapping class group will then contain these rotations, and will have a direct product group

structure:

$$\mathrm{MCG}(D_3) \approx \mathrm{PSL}_2(\mathbb{Z}) \times \mathbb{Z} \qquad (\partial D \text{ fixed pointwise}). \tag{2.4}$$

The factor $\times\mathbb{Z}$ counts the number of twists of the boundary.

If we fix ∂D setwise, then isotopies can freely rotate ∂D and remove any twist imposed by the diffeomorphism. This is the same as if the boundary was a puncture, so we obtain

$$\mathrm{MCG}(D_3) \approx \mathrm{PSL}_2(\mathbb{Z}) \qquad (\partial D \text{ fixed setwise}). \tag{2.5}$$

We shall usually assume that boundaries are fixed setwise, not pointwise, unless otherwise noted.

2.3 Dehn twists

Recall the two Dehn twists on the torus, Eq. (1.7): how do they act on the punctures of $S_{0,4}$? We use the same symbol for the mapping class and the linear diffeomorphism constructed from T_i, and regard the T_i as acting on $S_{0,4}$ via the quotient by the hyperellipetic involution, as described in Section 2.1. We write $\overline{T}_i = \pi \circ T_i \circ \pi^{-1}$ for the induced action on $S_{0,4}$. These maps are well-defined since the T_i commute with the hyperelliptic involution.

Clearly $T_1 p_0 = T_2 p_0 = p_0$, and it is easy to check

$$\begin{aligned} T_1(p_1) &= p_2, & T_2(p_1) &= p_1, \\ T_1(p_2) &= p_1, & T_2(p_2) &= p_3, \\ T_1(p_3) &= p_3, & T_2(p_3) &= p_2. \end{aligned}$$

The action of T_1 on the two closed curves from Figs. 2.4a–2.4b is shown in Figs. 2.6a–2.6b. The effect of T_1 on $S_{0,4}$ is to interchange p_1 and p_2 clockwise: the red curve is invariant,

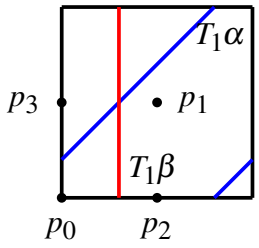

Fig. 2.6a The action of T_1 on the two curves from Fig. 2.4a.

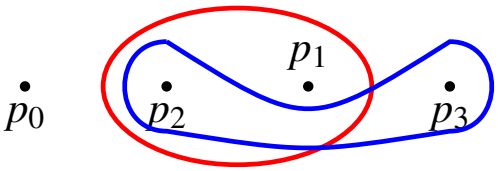

Fig. 2.6b The induced action of T_1, denoted $\overline{T}_1$, on the two curves from Fig. 2.4b.

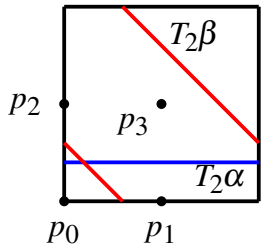

Fig. 2.7a The action of T_2 on the two curves from Fig. 2.4a.

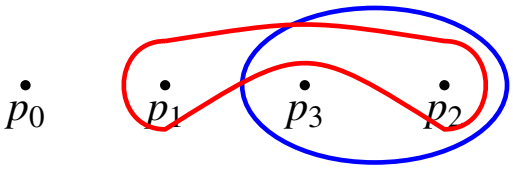

Fig. 2.7b The induced action of T_2, denoted $\overline{T}_2$, on the two curves from Fig. 2.4b.

since it contains p_1 and p_2; the blue curve is pulled and stretched below p_1 by the motion of p_2. Figures 2.7a–2.7b show the action of T_2 on the same curves: this time p_2 and p_3 are interchanged clockwise on $S_{0,4}$. We conclude that the two Dehn twists (1.7) on the torus map to 'Dehn half-twists' on $S_{0,4}$, which swap adjacent pairs of punctures clockwise. The inverse maps $\overline{T}_1^{-1}$ and $\overline{T}_2^{-1}$ swap punctures counterclockwise.

2.4 A stirring device example

A picture such as the rightmost frame in Fig. 2.5 can be interpreted as a circular container containing some fluid, with three movable stirring rods immersed [9, 20]. We neglect fluid motion perpendicular to the page, which is a good approximation when the fluid is thin, or stratified because of thermal effects. Now move the rods in some periodic fashion, so that they return to their initial positions as a set (they might have been permuted). The rods displace the fluid, which induces a diffeomorphism of D_3 to itself, which we call ϕ_1. This diffeomorphism could be computed numerically by solving a set of equations appropriate to the fluid: Stokes (slow viscous flow), Navier–Stokes, Oldroyd B (polymers), etc.

Now move the rods again in the same manner. This induces a diffeomorphism ϕ_2, which is not in general the same as ϕ_1: for instance, the system might be getting more and more turbulent at each period, such as when we stir cream into coffee. However, they must be isotopic, $\phi_1 \simeq \phi_2$, since the rods were moved the same way both times and the rods' periodic motion completely determines the mapping class (an element of $\mathrm{MCG}(D_3) \approx \mathrm{PSL}_2(\mathbb{Z})$, see Section 2.2). This means that we can use any of these maps, say ϕ_1, to represent the element of the mapping class $[\phi]$ induced by the rod motion.

Let us select a way of moving the rods by choosing an element of $\mathrm{MCG}(D_3)$. This corresponds to the isotopy class of a linear diffeomorphism $\phi : T^2 \to T^2$, which projects down to a diffeomorphism $\pi \circ \phi \circ \pi^{-1} = \bar{\phi} : D_3 \to D_3$ using the construction in Section 2.1.

This mapping is well-defined since linear diffeomorphisms on T^2 commute with the hyperelliptic involution ι. A good way of characterizing ϕ is to write it as a sequence of Dehn twists T_1 and T_2. Each of these Dehn twists maps to an interchange of punctures, as described in Section 2.3.

Let us illustrate this with a simple example, which is classical but was first discussed in the context of fluid dynamics by Boyland et al [9]. We start with an interchange of the first and second rods clockwise ($\overline{T}_1$), then interchange the second and third rods counterclockwise ($\overline{T}_2^{-1}$). The net result is the map $\bar{\phi} = \overline{T}_2^{-1}\overline{T}_1$. This lifts to a linear map on the torus $\phi = T_2^{-1}T_1$. The matrix representing this map is the Arnold cat map matrix (1.10), so we know the isotopy class of ϕ is Anosov (hyperbolic). The induced isotopy class of $\bar{\phi}$ on D_3 is called *pseudo-Anosov*.

What can we deduce about the action of $\bar{\phi}$? First, since ϕ is Anosov, any essential simple closed curve will asymptotically grow exponentially under repeated action of ϕ. But since essential simple closed curves on T^2 map to such curves on D_3 (see Figs. 2.4a–2.4b), we will observe exponential growth on D_3 as well. Thus, the fluid motion on D_3 carries at least the same 'complexity' as that on T^2.

Fig. 2.8 Experimental stirring device with three rods immmersed in a viscous fluid. Left: motion corresponding to $\overline{T}_2^{-1}\overline{T}_1$. Right: motion corresponding to $\overline{T}_2\overline{T}_1$. The black dye suggests how the fluid is displaced. (From Boyland et al [9].)

2.5 A taffy puller example

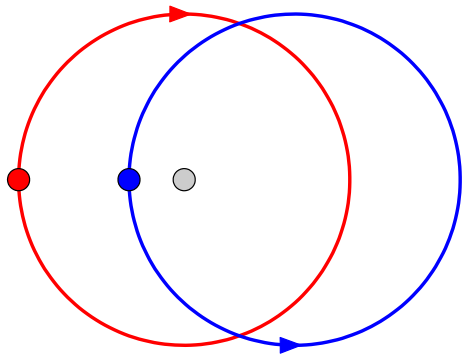

Fig. 2.9a The orbits of the rods in the taffy puller in Fig. 2.1b.

The motion of the rods in the taffy puller of Fig. 2.1b is shown in Fig. 2.9a. The two moving rods never cross each other's path. During their motion they only excchange position with the fixed rod. After a bit of starting at Fig. 2.9a we see that the rods induce the map $\overline{T}_2^{-1}\overline{T}_1^2\overline{T}_2^{-1}$. Here there is no ambient space: we simply imagine that the rods are puncturing a disk, which we use to measure the stretching motion. The linear diffeomorphism on T^2 corresponding to the rod motion is then

$$T_2^{-1}T_1^2T_2^{-1} = \begin{pmatrix} 3 & 4 \\ 2 & 3 \end{pmatrix} \tag{2.6}$$

with dilatation $\lambda = 3+2\sqrt{2}$. This gives us the factor by which the length of the taffy is multiplied at each full period. In that sense it is a measure of the effectiveness of the taffy puller.

Lecture 3
The braid group

Although it has been proved that every braid can be deformed into a similar normal form the writer is convinced that any attempt to carry this out on a living person would only lead to violent protests and discrimination against mathematics. He would therefore discourage such an experiment.
— Emil Artin, Theory of Braids

At the end of Lecture 2 we studied the motion of rods defining a 3-rod mixing device or taffy puller. We then associated these motions with the Dehn half-twists $\overline{T}_1$ and $\overline{T}_2$ that describe the clockwise interchange of punctures. If we want to study motions involving more rods, it is natural to introduce the *braid group* B_n for some fixed number of punctures n. This will allow us to encode much more complex motions.

3.1 Braids as particle dances

Artin [1] first introduced braids as collections of strings connecting points. He gave intuitive justification for several theorems, and in [2] provided rigorous proofs. There are several ways to define braids [5, 19, 6], but here we use the 'braids as particle dances' viewpoint, since it is intimately connected to dynamics. Consider n particles (or punctures) located at points in the Euclidean plane $\mathbb{E}^2$, which we regard as points in the complex plane $\mathbb{C}$. Now assume the points can move,

$$z(t) = (z_1(t), \ldots, z_n(t)), \qquad t \in I = [0,1], \tag{3.1}$$

such that

$$z_j(t) \neq z_k(t), \qquad j \neq k. \tag{3.2}$$

These last conditions mean that the particles never collide during their motion. Furthermore, assume that the points collectively return to their initial position:

$$\{z_1(0), \ldots, z_n(0)\} = \{z_1(1), \ldots, z_n(1)\} \qquad \text{(setwise)}. \tag{3.3}$$

The specification 'setwise' is important since it means the particles can be permuted amongst themselves. Such a vector of trajectories $z(t)$ defines a *braid*. If we embed $\mathbb{E}^2 \times I$

in $\mathbb{E}^3$ (what physicists call a space-time diagram), we obtain *geometric* or *physical* braids (Fig. 3.1a).

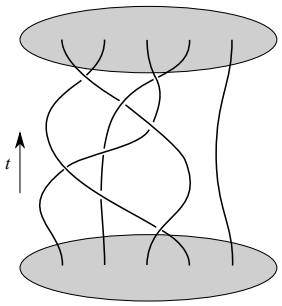

Fig. 3.1a A geometric braid.

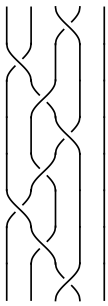

Fig. 3.1b The standard braid diagram corresponding to the braid on the left.

Intuitively, we should be able to deform a braid without really changing it, as long as the strings do not cross. This leads to the notion of equivalence of braids. Two braids are equivalent if there exists a homotopy from the first braid to the second, fixing the endpoints, such that at any time during the deformation the homotopy gives a valid braid. This defines an isotopy between the braids [5]. Since we will essentially be interested only in equivalence classes of braids under isotopy, and not in specific physical braids, from now on we drop 'equivalence class' and assume that equality of braids means equivalence under isotopy.

Now consider two braids α and α', with coordinate vectors $z(t)$ and $z'(t)$, sharing the same endpoints. We define the composition $\alpha'' = \alpha\alpha'$ as the braid with coordinate vector

$$z''(t) = \begin{cases} z(2t), & 0 \le t \le \frac{1}{2}; \\ z'(2t-1), & \frac{1}{2} \le t \le 1. \end{cases} \tag{3.4}$$

With this composition law, braids form a group B_n, the *Artin braid group on n strings*. It is easy to see that this composition law is associative. The identity braid has coordinate vector $z(t) = (z_1(0), \ldots, z_n(0))$. The inverse of α is the braid α^{-1} with coordinates $z(1-t)$. All the group axioms are thus satisfied.

3.2 Algebraic braids

It is common to redraw a braid as a *standard braid diagram*, where the braid is 'flattened' and where crossings occur one at a time in intervals of the same length (Fig. 3.1b). There

are various conventions for how to plot these, but here we read the 'time' parameter of the braid as flowing from bottom to top, in the same manner as a space-time diagram.

Figure 3.1b suggests that we rewrite a braid in terms of *standard generators*, denoted σ_i for $i = 1,\ldots,n$. The generator σ_i is the clockwise interchange of the ith string with the $(i+1)$th string, and σ_i^{-1} is the corresponding counterclockwise exchange. Any braid can be written as a product of these standard generators and their inverses. They satisfy the relations

$$\sigma_j\sigma_k = \sigma_k\sigma_j, \quad |j-k| > 1; \qquad \sigma_j\sigma_k\sigma_j = \sigma_k\sigma_j\sigma_k, \quad |j-k| = 1. \tag{3.5}$$

These relations are depicted in Figs. 3.2a–3.2b. The first type of relation says that gener-

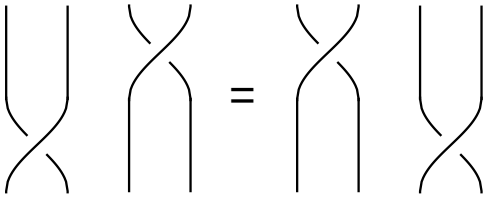

Fig. 3.2a Relation $\sigma_j\sigma_k = \sigma_k\sigma_j$.

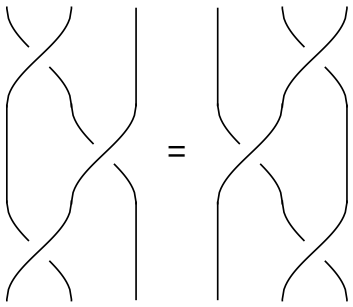

Fig. 3.2b Relation $\sigma_j\sigma_k\sigma_j = \sigma_k\sigma_j\sigma_k$.

ators that don't share a string commute. The second is less obvious, and involves triplets of adjacent strings; staring at Fig. 3.2b for a while should convince the reader that the relations holds. The latter are the key defining relations for the braid group, and are often referred as the *braid relations* when they occurs in other groups. Artin [2] proved the somewhat surprising fact that (3.5) are the only nontrivial relations for the braid group B_n. A braid given by a product of generators obeying (3.5) is called an *algebraic braid*, since it is not necessarily associated with a physical braid.

The (positive) *half-twist* is the element $\Delta_n \in B_n$ defined as

$$\Delta_n = (\sigma_{n-1}\sigma_{n-2}\cdots\sigma_1)(\sigma_{n-1}\sigma_{n-2}\cdots\sigma_2)\cdots(\sigma_{n-1}\sigma_{n-2})(\sigma_{n-1}). \tag{3.6}$$

As its name implies, the half-twist consists of grabbing the whole braid as a ribbon and giving it a 180° twist. Note that $\sigma_j\Delta_n = \Delta_n\sigma_{n-j}$, which then implies that $\sigma_j\Delta_n^2 = \Delta_n^2\sigma_j$, i.e., Δ_n^2 commutes with each generator. This means that Δ_n^2 is in the *center* of B_n — the subgroup consisting of elements that commute with every element of B_n. In fact the center of the braid group is generated by powers of Δ_n^2, which is called the (positive) *full-twist*.

The braid group is most important to us for its connection to mapping class groups of punctured disks and spheres. First, observe that

$$\mathrm{MCG}(D_n) \approx B_n, \qquad (\partial D \text{ fixed pointwise}); \tag{3.7}$$

that is, the mapping class group of disk with n puntures is isomorphic to B_n. Here the boundary of the disk is fixed pointwise by the diffeomorphisms and isotopies. Comparing this to (2.4), we see that $B_n \approx \mathrm{PSL}_2(\mathbb{Z}) \times \mathbb{Z}$. The factor $\times\mathbb{Z}$ counts the number of full-twists.

If the boundary of the disk is fixed as a set rather than pointwise, we have

$$\mathrm{MCG}(D_n) \approx B_n/\langle \Delta_n^2 \rangle, \qquad (\partial D \text{ fixed setwise}). \tag{3.8}$$

In this case the diffeomorphisms can 'rotate' the boundary, so the isotopy can unwind any net full-twist. Hence the quotient by the center that appears in (3.8). Comparing this to (2.5), we see that $B_n/\langle \Delta_n^2 \rangle \approx \mathrm{PSL}_2(\mathbb{Z})$.

3.3 Artin's representation

Consider the generating set of loops for the fundamental group $\pi_1(D_n, x_0)$ in Fig. 3.3. Unlike for the torus case (Lecture 1), the set $\langle e_1, \cdots, e_n \rangle$ form a *free group*, that is,

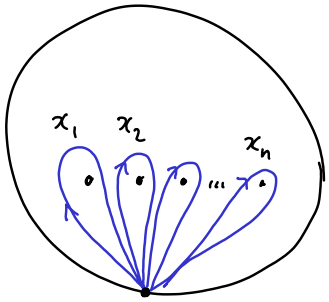

Fig. 3.3 A generating set of loops for $\pi_1(D_n, x_0)$.

$\pi_1(D_n, x_0) \approx F_n$. The free group with n generators F_n is a finitely-generated group where the generators obey no nontrivial relations. (Unlike the braid group B_n, whose generators satisfy the relations (3.5).) Free groups have the very useful property that two words (i.e., product of generators) are equal if and only if they are equal generator-by-generator (also called lexicographical equality).

Recall that the generator $\sigma_i \in B_n$ exchanges the position of the ith and $(i+1)$th punctures. Since elements of B_n correspond to elements of $\mathrm{MCG}(D_n)$ (see (3.7)–(3.8)), there is

a canonical diffeomorphism corresponding to σ_i whose induced action on loops is pictured in Fig. 3.4. The loops e_j for $j \neq i, i+1$ are untouched.

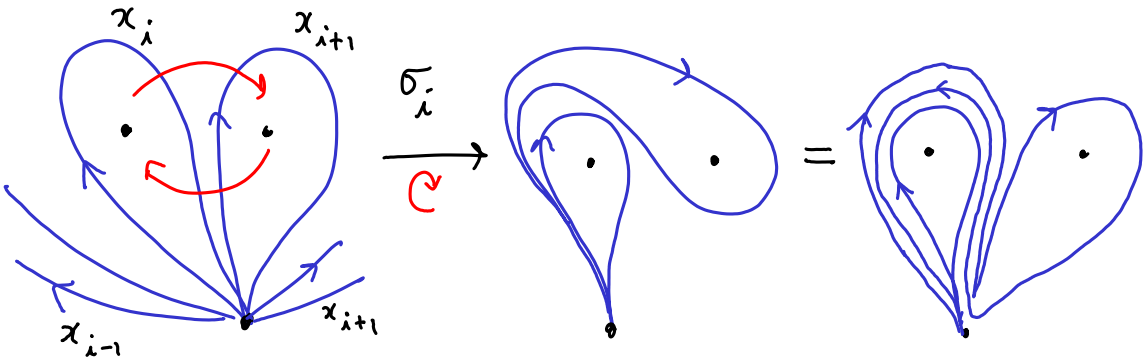

Fig. 3.4 The diffeomorphism corresponding to σ_1 acts on the generating loops.

We write $\sigma_{i*} : \pi_1(D_n, x_0) \to \pi_1(D_n, x_0)$ for the induced action of σ_i. In mathematical terms, $\sigma_{i*} \in \mathrm{Aut}(F_n)$, the group of *automorphisms* of the free group. The rightmost picture in Fig 3.4 shows us how to express the result of σ_{i*} in terms of the initial loops:

$$\begin{aligned} e_i\,\sigma_{i*} &= e_i e_{i+1} e_i^{-1}\,, \\ e_{i+1}\,\sigma_{i*} &= e_i\,, \\ e_j\,\sigma_{i*} &= e_j\,, \qquad j \neq i, i+1. \end{aligned} \tag{3.9}$$

Note that we use a *right action*, since in our convention braid words are read from left to right, and so generators must act to their left to maintain the correct order. Artin showed that mapping each σ_i to σ_{i*} defines a faithful representation of the braid group B_n. For instance,

$$\begin{aligned} e_i\,\sigma_{i*}\sigma_{i+1*}\sigma_{i*} &= (e_i e_{i+1} e_i^{-1})\sigma_{i+1*}\sigma_{i*} \\ &= (e_i\sigma_{i+1*})(e_{i+1}\sigma_{i+1*})(e_i^{-1}\sigma_{i+1*})\sigma_{i*} \\ &= (e_i\, e_{i+1} e_{i+2} e_{i+1}^{-1}\, e_i^{-1})\sigma_{i*} \\ &= e_i e_{i+1} e_i^{-1}\, e_i\, e_{i+2}\, e_i^{-1}\, e_i e_{i+1}^{-1} e_i^{-1} = e_i e_{i+1}\, e_{i+2}\, e_{i+1}^{-1} e_i^{-1}\,, \end{aligned} \tag{3.10}$$

and

$$\begin{aligned} e_i\,\sigma_{i+1*}\sigma_{i*}\sigma_{i+1*} &= e_i\,\sigma_{i*}\sigma_{i+1*} \\ &= e_i e_{i+1} e_i^{-1}\,\sigma_{i+1*} \\ &= (e_i\sigma_{i+1*})(e_{i+1}\sigma_{i+1*})(e_i^{-1}\sigma_{i+1*}) = e_i\, e_{i+1} e_{i+2} e_{i+1}^{-1}\, e_i^{-1}. \end{aligned} \tag{3.11}$$

We see that (3.10) is the same as (3.11). We can repeat this for the action on the other generating loops, and we find that $\sigma_{i*}\sigma_{i+1*}\sigma_{i*} = \sigma_{i+1*}\sigma_{i*}\sigma_{i+1*}$. We can verify that all the braid relations (3.5) are thus satisfied, and we have a representation of B_n. This is not a linear representation in terms of matrices, since it acts on a (non-Abelian) free group.

We can obtain a representation in terms of matrices by *Abelianization*: take the action (3.9) and instead of the e_i being in F_n, assume that they commute. This is the *group ring* with integer coefficients $\mathbb{Z}F_n$, and we use the same symbol for $e_i \in \mathbb{Z}F_n$. The e_i form a basis for $H_1(D_n,\mathbb{Z})$, the first homology group of D_3 with integer coefficients. The only nontrivial actions are $e_i\,\sigma_{i*} = e_{i+1}$ and $e_{i+1}\,\sigma_{i*} = e_i$. This is simply the permutation of e_i and e_{i+1}: we have recoved a representation for the symmetric group on n symbols in terms of elementary permutation matrices. This is not very useful: it captures nothing of the special character of the braid group.

Let us summarize: the representation of B_n in terms of automorphisms $\mathrm{Aut}(F_n)$ is faithful, meaning that it is the identity if and only if the corresponding braid is the identity. However, such a representation is hard to manipulate, as is obvious from (3.10)–(3.11). Abelianizing gives us matrices acting on $H_1(D,\mathbb{Z})$, which are very easy to multiply together. However, the resulting representation can only characterize the permutation induced by a braid, so the representation is extremely unfaithful and is a poor characterization of the braid group B_n. Is there a way to have our cake and eat it too: to have a representation in terms of matrices that is also faithful? We shall see in the next section that this is (almost) possible.

3.4 The Burau representation

Figure 3.5b shows the disk D_3 with a generating set for its fundamental group. We have added three *branch cuts*: whenever a loop crosses such a cut, it now exists on a new copy (or *deck*) of the punctured disk. This collection of decks is called a *branched cover* of D_n, denoted $\widetilde{D}_n$. This construction is shown schematically in Fig. 3.5b. Note that the three cuts share the same new copy of the disk (unlike the *universal cover*, where each cut would get its own new copy). It is easy to see how to generalize this construction to more punctures. We call this new space a $\mathbb{Z}$*-cover* of the disk D_n, since each deck can be indexed by a single number in $\mathbb{Z}$. (See the lectures by Boyland and Franks [8].)

When lifted to the $\mathbb{Z}$ cover, a loop e_i is no longer closed, and becomes an arc. It joins the basepoint x_0 to another point tx_0. Here t is a *deck transformation* which moves a point from one deck of the cover to the next one. In general, t^m with $m \in \mathbb{Z}$ moves us from the fundamental domain (the original disk) to the mth deck. We write $\tilde{e}_i$ for the unique *lift* of e_i to the $\mathbb{Z}$-cover. Then $t^m\tilde{e}_i$ is a curve joining the basepoint in deck $m-1$ to the basepoint in

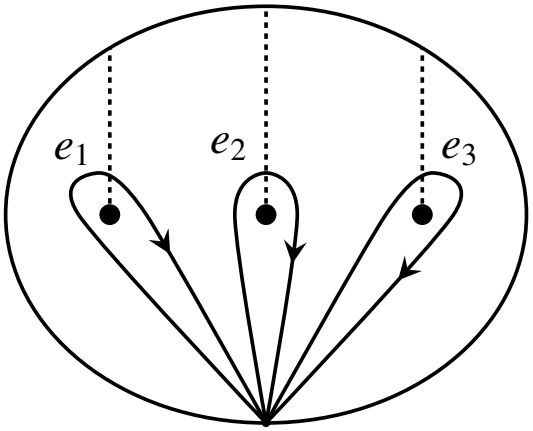

Fig. 3.5a The disk with 3 punctures and a generating set for $\pi_1(D_3, x_0)$. The dashed lines are branch cuts.

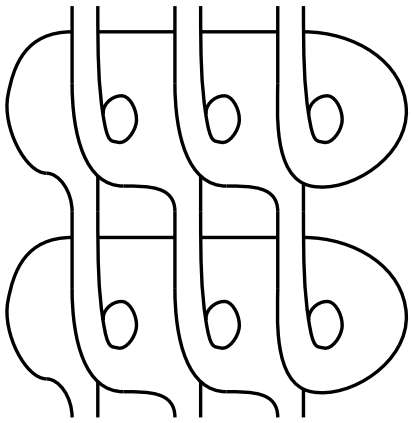

Fig. 3.5b Branched cover $\widetilde{D}_3$ corresponding to the cuts on the left. Two decks are shown, but the cover can be infinite in each direction, or the decks at each end can eventually join up.

deck m. In Fig. 3.6 (left) we represent the curves and the different decks of the $\mathbb{Z}$-cover as a graph (the *1-skeleton* of the $\mathbb{Z}$-cover).

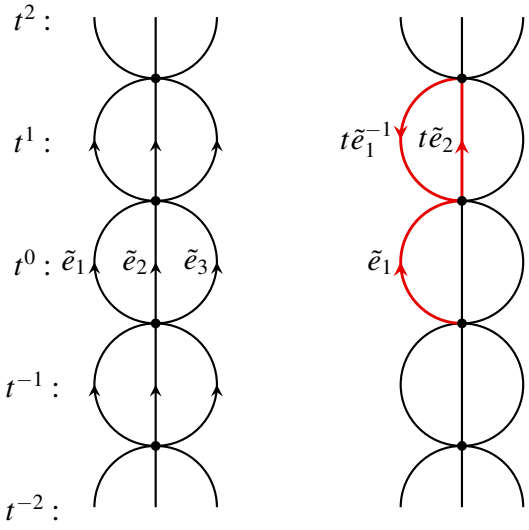

Fig. 3.6 Left: the 1-skeleton of the $\mathbb{Z}$-cover $\widetilde{D}_3$. Right: the action of σ_{1*} on the three generating curves.

Now consider the same action as (3.9) induced by σ_{i*} on the $\mathbb{Z}$-cover D_n. We have

$$\begin{aligned} \tilde{e}_i\,\sigma_{i*} &= \tilde{e}_i(t\tilde{e}_{i+1})(t\tilde{e}_i^{-1})\,, \\ \tilde{e}_{i+1}\,\sigma_{i*} &= \tilde{e}_i\,, \\ \tilde{e}_j\,\sigma_{i*} &= \tilde{e}_j\,, \qquad j \neq i, i+1. \end{aligned} \tag{3.12}$$

The first line of this action is depicted on the 1-skeleton in Fig. 3.6 (right) for $i = 1$ and $n = 3$. An easy way to see how this works is to remember that the curves are continuous, so curve $\tilde{e}_1$ must be followed by a curve that begins at tx_0, in this case $t\tilde{e}_2$. The final segment $t\tilde{e}_1^{-1}$ actually begin at t^2x_0, since it is the reverse of $t\tilde{e}_1$.

What have we gained by lifting to the $\mathbb{Z}$-cover? So far, not much. But the trick is that now we Abelianize (3.12); the first line then reads

$$\tilde{e}_i\,\sigma_{i*} = \tilde{e}_i + t\tilde{e}_{i+1} - t\tilde{e}_i = (1-t)\,\tilde{e}_i + t\tilde{e}_{i+1} \tag{3.13}$$

where following convention we now use addition and subtraction to write the Abelian composition law. Because of the parameter t, there is no cancellation in (3.13) and we do not retain only the permutations. The representation of B_n is now in terms of an action on $H_1(\widetilde{D}_n, \mathbb{Z})$, the first homology group of $\widetilde{D}_n$ with integer coefficients. The elements of the homology group have the form $t^m\tilde{e}_i$.

Using (3.13) and the other rows of (3.12) (which are unchanged by Abelianization), we can thus represent the action of σ_{i*} as a matrix, which gives the *Burau representation* of B_n [10, 5, 6]:

$$\sigma_i \;\mapsto\; I_{i-1} \oplus \begin{pmatrix} 1-t & t \\ 1 & 0 \end{pmatrix} \oplus I_{n-i-1} \tag{3.14}$$

where I_m denotes an m by m identity matrix. The Burau representation maps B_n to $\mathrm{GL}_n(\mathbb{Z}[t,t^{-1}])$, the general linear group with entries that are integer polynomials in t and t^{-1}.

The matrices of the Burau representation have an invariant vector:

$$u_n = \tilde{e}_1 + t\tilde{e}_2 + \cdots + t^{n-1}\tilde{e}_n \tag{3.15}$$

which corresponds to the lift of a loop around all the punctures, $e_1e_2\cdots e_n$. Such a loop is clearly invariant under any diffeomorphism induced by the σ_is. If we now change to a basis $u_i = \tilde{e}_i - \tilde{e}_{i+1}$, $1 \le i \le n-1$, and u_n as above, the matrices of the representation block up, with a 1×1 block in the bottom-right slot. We may thus drop the last coordinate to obtain the *reduced Burau representation*, which maps B_n to $\mathrm{GL}_{n-1}(\mathbb{Z}[t,t^{-1}])$:

$$[\sigma_i] = I_{i-2} \oplus \begin{pmatrix} 1 & t & 0 \\ 0 & -t & 0 \\ 0 & 1 & 1 \end{pmatrix} \oplus I_{n-i-2}\,. \tag{3.16}$$

Here for $i = 1$ we interpret I_{i-2} as deleting the first row and column of the 3×3 block, and for $i = n-1$ we interpret I_{n-i-2} as deleting the last row and column of that block. The $-t$ entry in (3.16) is always in the diagonal (i,i) position. Thus, for $n = 3$ we have

$$[\sigma_1] = \begin{pmatrix} -t & 0 \\ 1 & 1 \end{pmatrix}, \qquad [\sigma_2] = \begin{pmatrix} 1 & t \\ 0 & -t \end{pmatrix}, \tag{3.17}$$

and for $n = 4$

$$[\sigma_1] = \begin{pmatrix} -t & 0 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \qquad [\sigma_2] = \begin{pmatrix} 1 & t & 0 \\ 0 & -t & 0 \\ 0 & 1 & 1 \end{pmatrix}, \qquad [\sigma_3] = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & t \\ 0 & 0 & -t \end{pmatrix}. \tag{3.18}$$

From now on when we refer to the Burau representation we shall usually mean the reduced version.

What is the role of the parameter t? So far we have treated it as simply a symbol. But it can be useful to assign an actual complex value to t. For $t = 1$, we can see from (3.14) that the (unreduced) Burau representation just becomes a representation of the symmetric group, so that value should be avoided as it loses too much information about B_n. If we take t to be a root of unity with $t^M = 1$, then the $\mathbb{Z}$-cover becomes a finite cover of degree M (a $\mathbb{Z}_M$-cover).

For $t = -1$ we get a double-cover ($t^2 = 1$). For D_3 in particular observe that with $t = -1$ the matrices (3.17) are identical to the matrices T_1 and T_2 of the Dehn twists on the torus given by (1.7). This is because the double cover reverses the quotient of the torus by the hyperelliptic involution that we took in Lecture 1. Thus, for $t = -1$ and $n = 3$ the reduced Burau matrices generate $\mathrm{SL}_2(\mathbb{Z})$.

Is the Burau representation faithful? Recall that the representation is faithful if the kernel of the map $B_n \to \mathrm{GL}_{n-1}(\mathbb{Z}[t,t^{-1}])$ is trivial. This means that the only braid that maps to the identity matrix is the identity braid. It has long been known that the representation is faithul for $n \leq 3$. Then Moody [17] showed it is unfaithful for $n \geq 9$. Long and Paton [15] improved this to $n \geq 6$. Finally, Bigelow [3] showed that it is unfaithful for $n = 5$. This leaves the case $n = 4$, which is still open. However, Bigelow [4] showed that another matrix representation, the Lawrence–Krammer representation, is faithful for all n.

The lack of faithfulness of the Burau representation should not trouble us too much: it does not affect the central purpose for which we'd like to use the Burau representation — to get lower bounds on the dilatation of braids. We shall see how to do this in Section 3.5.

3.5 Lower bounds on topological entropy

In Section 1.3.3 we introduced the dilatation λ of a mapping class of the torus, with $h = \log\lambda$ the *topological entropy* of the mapping class. In fact the topological entropy $h(\phi)$ can be defined for any diffeomorphism ϕ on a surface S. A useful way to think of topological entropy comes from a bound due to Bowen [7]: he proved that $\log\lambda$ for a surface S is bounded below by the *growth* on $\pi_1(S)$. (We omit the basepoint for simplicity.)

For a diffeomorphism ϕ, the growth on $\pi_1(S)$ is a measure of how fast elements of $\pi_1(S)$ grow under repeated action of ϕ_*. We define the *reduced word length* $\|\gamma\|$ of an element $\gamma \in \pi_1(S)$ as the minimum number of generators e_i needed to write γ. When $\pi_1(S)$ is a free group this is easy to compute: all we have to do is count generators after cancelling adjacent inverses. The *growth* of $\pi_1(S)$ under the action of ϕ_* is

$$\mathrm{GR}(\phi_*) = \sup_{\gamma\in\pi_1(S)} \limsup_{m\to\infty} \|\phi_*^m\gamma\|^{1/m}. \tag{3.19}$$

Bowen [7], building on a theorem of Manning [16], showed that

$$h(\phi) \geq \log \mathrm{GR}(\phi_*). \tag{3.20}$$

The difficulty with Bowen's bound (3.20) is that the growth on $\pi_1(S)$ given by (3.19) is rather hard to compute. Many of the techniques we will introduce are essentially different methods of computing $\mathrm{GR}(\phi_*)$.

For the case where $S = D_n$, the disk with n punctures, one of the most direct methods is not an exact calculation, but rather gives another bound. Kolev [14] proved directly a theorem that is also a corollary of Fried [13]. Recall that the mapping class group of D_n is isomorphic to the braid group ((3.7) and (3.8)). Given a mapping class represented by the braid $\alpha \in B_n$, we have [14]

$$\mathrm{GR}(\phi_*) \geq \sup_{|t|=1} \mathrm{spr}([\alpha](t)) \tag{3.21}$$

where $[\alpha](t)$ is the Burau representation of α from Section 3.4, with parameter $t \in \mathbb{C}$, and $\mathrm{spr}(A)$ is the *spectral radius* of the matrix A — the largest magnitude of its eigenvalues.

References

[1] Artin E (1925) Theorie der zöpfe. Abh Math Semin Univ Hambg 4(1):47–72

[2] Artin E (1947) Theory of braids. Ann Math 48(1):101–126

[3] Bigelow SJ (1999) The Burau representation is not faithful for $n = 5$. Geom Topol 3:397–404

[4] Bigelow SJ (2001) Braid groups are linear. J Amer Math Soc 14(2):471–486

[5] Birman JS (1975) Braids, Links and Mapping Class Groups. No. 82 in Annals of Mathematics Studies, Princeton University Press, Princeton, NJ

[6] Birman JS, Brendle TE (2005) Braids: A survey. In: Menasco W, Thistlethwaite M (eds) Handbook of Knot Theory, Elsevier, Amsterdam, pp 19–104, available at `http://arXiv.org/abs/math.GT/0409205`

[7] Bowen R (1978) Entropy and the fundamental group. In: Structure of Attractors, Springer, New York, Lecture Notes in Math., vol 668, pp 21–29

[8] Boyland PL, Franks JM (1989) Lectures on dynamics of surface homeomorphisms. University of Warwick Preprint, notes by C. Carroll, J. Guaschi and T. Hall

[9] Boyland PL, Aref H, Stremler MA (2000) Topological fluid mechanics of stirring. J Fluid Mech 403:277–304

[10] Burau W (1936) Über Zopfgruppen und gleichsinnig verdrilte Verkettungen. Abh Math Semin Univ Hambg 11:171–178

[11] Farb B, Margalit D (2011) A Primer on Mapping Class Groups. Princeton University Press, Princeton, NJ

[12] Finn MD, Thiffeault JL (2011) Topological optimization of rod-stirring devices. SIAM Rev 53(4):723–743, http://arXiv.org/abs/1004.0639

[13] Fried D (1986) Entropy and twisted cohomology. Topology 25(4):455–470

[14] Kolev B (1989) Entropie topologique et représentation de Burau. C R Acad Sci Sér I 309(13):835–838, english translation at `http://arxiv.org/abs/math.DS/0304105`.

[15] Long DD, Paton M (1993) The Burau representation is not faithful for $n \geq 6$. Topology 32:439–447

[16] Manning A (1978) Topological entropy and the first homology group. In: Dynamical Systems — Warwick 1974, Springer, Berlin, Lecture Notes in Mathematics, vol 468, pp 185–190

[17] Moody JA (1991) The Burau representation of the braid group B_n is unfaithful for large n. Bull Am Math Soc 25:379–384

[18] Munkres JR (2000) Topology, 2nd edn. Prentice-Hall, Upper Saddle River, NJ

[19] Rolfsen D (2003) New developments in the theory of Artin's braid groups. Topology Appl 127(1-2):77–90, DOI 10.1016/S0166-8641(02)00054-8

[20] Thiffeault JL, Finn MD (2006) Topology, braids, and mixing in fluids. Philos Trans Royal Soc Lond A 364(1849):3251–3266, DOI 10.1098/rsta.2006.1899