AN INTRODUCTORY COURSE ON BRAIDS

FLORIAN DELOUP

Abstract. TBW

Contents

1. Introduction

The braid group on n strands – denoted B_n – was introduced by Emil Artin in 1925. It is one of the most ubiquitous group in mathematics: it is part of a classical family of groups (Artin–Tits groups), it plays a central role in low-dimensional topology (dim 2 and 3), in the construction and understanding of invariants of knots and links, it is related to the representation theory of quantum groups (quasi-Hopf algebras), the monodromy techniques in algebraic geometry and lies at the heart of Grothendieck–Drinfeld's outstanding program of understanding $\mathrm{Gal}(\overline{\mathbb{Q}} : \mathbb{Q})$ in geometric terms. More recently, the braid group has become the key ingredient in F. Cohen's program in the computation of homotopy spheres.

Here is the definition: the group B_n is defined by the presentation

$$\sigma_1, \ldots, \sigma_{n-1} \mid \sigma_i\sigma_j = \sigma_j\sigma_i \ (|i-j| \geq 2),\ \sigma_i\sigma_{i+1}\sigma_i = \sigma_{i+1}\sigma_i\sigma_{i+1} \ (i = 1, 2, \ldots, n-2).$$

There is a braid group on infinitely many strands,denoted B_∞, defined by a presentation with countably infinitely many generators $\sigma_1, \ldots$ and same relations.

There are two kinds of relations: commutativity and triangle relations. Except for B_2, which is easily seen to be isomorphic to $\mathbb{Z}$, all braid groups are noncommutative groups.

We now proceed to other, equivalent, definitions.

1.1. **First geometric definition.** Denote by D the unit disk in $\mathbb{C}$ and D_n the unit disk with n distinguished (fixed) points. (puncture points). An *n-strand geometric braid* is the image of an embedding b of n disjoint copies $[0_1, 1_1], \ldots [0_n, 1_n]$ of the unit interval $[0,1]$ into the cylinder $[0,1] \times D$ with the following two properties:

(i) $\{b(0_1), \ldots, b(0_n)\}$ is the set of puncture points of $D_n = 0 \times D$ and $\{b(1_1), \ldots, b(1_n)\}$ is the set of puncture points of $D_n = 1 \times D$.

(ii) For each t in $[0_j, 1_j]$, $b(t) \in \{t\} \times D$.

The image of each interval is called a *strand* of the braid. Two geometric braids b, b' are *isotopic* if there is a smooth family of geometric braids $(b_t)_{t\in[0,1]}$ such that$b_0 = b$ and $b_1 = b'$. Key points: 1) The endpoints are fixed throughout; 2) The map $t \mapsto b(t)_1$ is strictly increasing.

A couple of easy observations:

1. A geometric braid induces a permutation of the puncture points.
2. Isotopic braids induce equal permutations.

We define a *braid* to be an isotopy class of geometric braids.

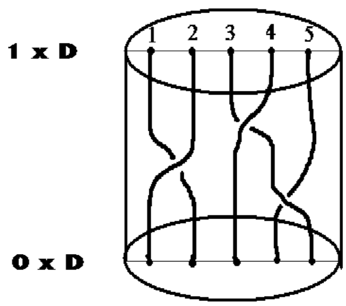

FIGURE 1.1. A geometric braid with 5 strands.

First piece of algebra (see Figures 1.1 and 1.1): the set of braids (with a fixed number of strands) has a group structure given by concatenation.

- Given two braids b, b', squeeze the image of b into the cylinder $[0, 1/2] \times D$ and the image of b' into the cylinder $[1/2, 1] \times D$; obtain a new, well-defined, braid $b \cdot b' = bb'$.
- There is a trivial element: it is represented by a geometric braid whose strands are straight segements, not intertwining each other.

- The inverse of a braid is induced by its reflection in the disk $\{1/2\} \times D$.

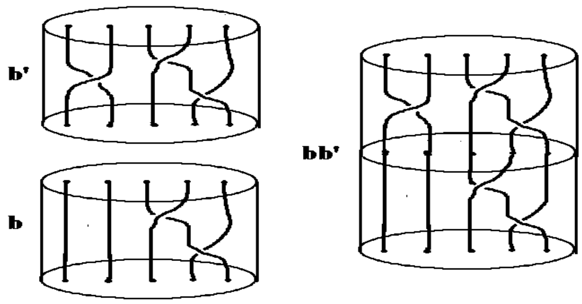

FIGURE 1.2. Multiplication of braids.

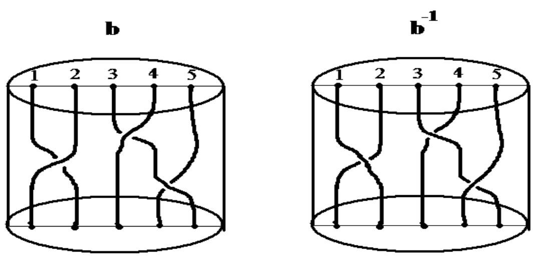

FIGURE 1.3. The inverse of a braid.

It is easy to see that n-strand braids are generated by $n-1$ elements $s_1, \dots, s_{n-1}$. For each j, s_j is represented by a geometric braid with $n-2$ straight strands not intertwining with any other strand, and the j-th and $(j+1)$-th strands intertwin exactly once. There is a well-known orientation convention to specify to make this completely well-defined. The associated permutation is the transpostion $(j, j+1)$ which exchanges j and $j+1$.

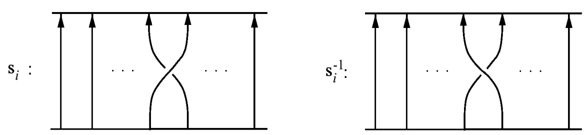

FIGURE 1.4. The Artin generators of the braid group.

Let us denote by $\mathfrak{B}_n$ the group of braids with n strands according to the geometric definition.

Theorem 1.1. *The groups B_n and $\mathfrak{B}_n$ are isomorphic.*

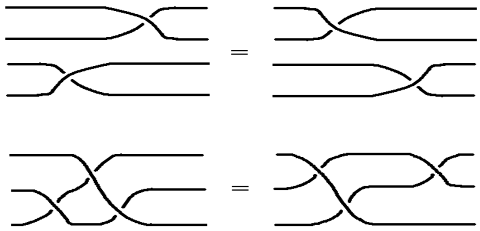

FIGURE 1.5. The braid relations viewed geometrically.

Proof. There is a well-defined homomorphism sending σ_j to s_j : the relations are easy to check with pictures (they are locally represented on figure 1.1). The relation $s_i s_{i+1} s_i = s_{i+1} s_i s_{i+1}$ is the well-known Reidemeister III relation in knot theory. Surjectivity: any geometric braid can be deformed into one in which a side view has only finitely many crossings, all of which are transverse. Injectivity: The two types of relations are sufficient to transform into each other any two braid diagrams representing isotopic braids. To see this in a purely combinatorial way, it is convenient to turn to the PL (Piecewise-linear) category (this can also be done in the locally analytic category). First step: each isotopy can be realized by means of isotopies which are affine: only one segment is non-trivially involved in the isotopy and at the end of the isotopy, it is replaced by a broken segment, with the endpoints of the original segment remaining fixed. Second step (exercise): look at the possible configurations in triangles. Only the two relations mentionned above can appear, all other relations can be (further) simplified. ■

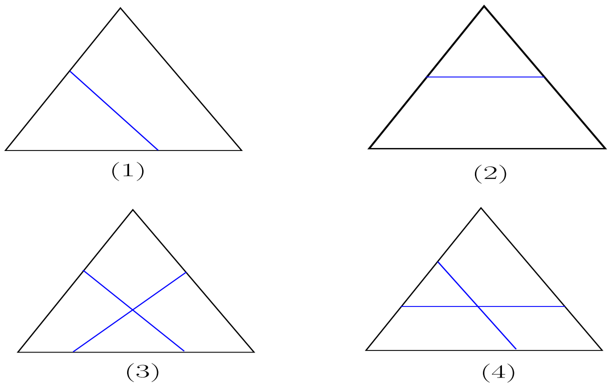

FIGURE 1.6. Configurations.

Solution of the exercise. Inside the triangle, we can assume that there is at most one double point (otherwise, we can decompose the isotopy further as a composition of affine isotopies). Up to symmetry, there are 4 configurations to consider. See figure 1.1. Configuration (1) yields an isotopy which either does not modify the word of the geometric braid or is a relation of type $\sigma_i\sigma_j = \sigma_j\sigma_i$ (for $|i-j|>1$). Configuration (2) describes an isotopy which involves a relation of type $\sigma_i\sigma_i^{-1} = e$ or $\sigma_i^{-1}\sigma_i = e$. Configuration (3) describes an isotopy which involves a relation of type $\sigma_i\sigma_{i+1}\sigma_i = \sigma_{i+1}\sigma_i\sigma_{i+1}$. Configuration (4) describes an isotopy which is the composition of two relations, one of each type. ■

Remark 1.2. The canonical projection $B_n \to S_n$ can be regarded as the composition of the isomorphism above $B_n \to \mathfrak{B}$ with the map $\mathfrak{B}_n \to S_n$ sending each braid s_i to its associated transposition (exchanging i and $i+1$).

Remark 1.3. The braid group, using this geometric definition, can be generalized as follows. For any surface Σ, compact, with boundary, not necessarily orientable, marked with n points, we can define $B_n(\Sigma)$ as the group of n-strand braids in $[0,1]\times\Sigma$.

1.2. **Mapping class group definition.** Let Σ be an oriented compact surface with or without boundary and $\mathcal{P}$ a set of distinguished points (punctures) lying in the interior of Σ. The *mapping class group* $\mathrm{MCG}(\Sigma,\mathcal{P})$ is the group of all isotopy classes of orientation-preserving homeomorphisms $\varphi : \Sigma \to \Sigma$ such that $\varphi|_{\partial\Sigma} = \mathrm{id}_{\partial\Sigma}$ and $\varphi(\mathcal{P}) = \mathcal{P}$.

Any self-homeomorphism of Σ sending punctures to punctures and preserving pointwise $\partial\Sigma$ represents an element in $\mathrm{MCG}(\Sigma,\mathcal{P})$. Punctures may be permuted by φ. The notion of isotopy is the obvious one. The group structure is induced by composition of homeomorphisms. It is known that all mapping class groups have *finite* presentations.

Theorem 1.4. *The groups* $\mathrm{MCG}(D_n)$ *and* B_n *are isomorphic.*

Proof. Let $\varphi : D_n \to D_n$ represent an element of $\mathrm{MCG}(D_n)$. Use Alexander's trick: any self-homeomorphism of a disk which fixes the boundary is isotopic to the identity through self-homeomorphisms fixing the boundary. Take such an isotopy $(\varphi_t)_{t\in[0,1]} : [0,1]\times D_n \to D_n$ between $\varphi(x) = \Phi(1,x)$ and id$=\Phi(0,x)$. Consider the map b from $[0,1]\times\mathcal{P}$ to $[0,1]\times D$, defined by

$$b(t,x_j) = (t,\varphi_t(x_j)),\quad t\in[0,1],\quad x_j\in\mathcal{P}.$$

It is an embedding of n disjoint unit intervals into the cylinder $[0,1]\times D$ clearly satisfying the defining properties of a geometrical braid. This construction defines a map

$$\mathrm{MCG}(D_n) \to B_n$$

which is easily seen to be a homomorphism. Conversely, let β be a geometric braid sitting in $[0,1]\times D$. It can be viewed as the graph of a motion of n pairwise distinct points moving in D. Extend the motion to a family of self-homeomorphisms of D starting at the identity, fixing the boundary. This yields an isotopy whose end map is the required self-homeomorphism. It is clear that these two constructions are inverse of each other. ■

We now define particular elements of the mapping class groups. Identify S^1 and $\mathbb{R}/\mathbb{Z}$. The *basic Dehn twist* τ is the map $[0,1] \times S^1 \to [0,1] \times S^1$ defined by

$$(t,x) \mapsto (t, t+x).$$

Observe that $\tau(0,x) = (0,x)$ and $\tau(1,x) = (1,x)$. In other words, τ preserves the boundary pointwise. Let now α be a simple closed curve in Σ. Choose a parametrization $h : S^1 \to \alpha$. Any such parametrization extends to a homeomorphism $h : [0,1] \times S^1 \to N(\alpha)$, where $N(\alpha)$ is a closed tubular neighborhood of α in Σ and α is the image of $1/2 \times S^1$. The orientation of $[0,1] \times S^1$ should match that of $N(\alpha)$ (induced by the orientation of Σ). The *Dehn twist* τ_α *around* α is the self-homeomorphism $\Sigma \to \Sigma$ defined by

$$\tau_\alpha(x) = \begin{cases} h \circ \tau \circ h^{-1}(x) & \text{if } x \in N(\alpha); \\ x & \text{if } x \in \Sigma - N(\alpha). \end{cases}$$

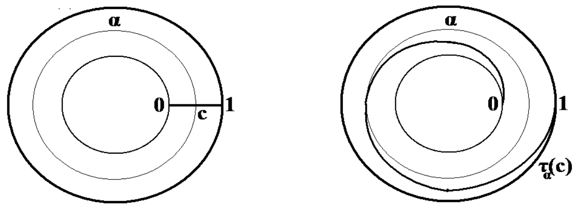

FIGURE 1.7. The basic Dehn twist and its effect on a (portion of a) curve.

The Dehn twist τ_α does not depend on the orientation of α. The element in $\mathrm{MCG}(\Sigma, \mathcal{P})$ corresponding to τ_α does not depend on the choice of the neighborhood of α.

There is a similar construction for an arc α whose endpoints p, q lie in the set $\mathcal{P}$ of punctures. Assume that the interior of α lie in $\Sigma - \mathcal{P}$. Fix a connected neighborhood $N(\alpha)$ of α in Σ:it is homeomorphic to a disk D. For simplicity, we forget about the parametrization. We may assume that $p = 1/4$ and $q = -1/4$ and that α is identified with the segment $[-1/4, 1/4]$. The annulus $A = \{z \in D \mid 1/2 \leq |z| \leq 1\}$ is identified to $[0,1] \times S^1$. The *basic Dehn half-twist* τ is defined by

$$\tau(t,x) = \begin{cases} (t,x) & \text{if } (t,x) \notin D \\ (t, \frac{t}{2} + x) & \text{if } (t,x) \in A = [0,1] \times S^1 \\ (t,-x) & \text{if } (t,x) \in D - A. \end{cases}$$

In other words, τ is the homeomorphism obtained from the isotopy of the identity map to the rotation about the midpoint of (x,y) by the angle π. *The Dehn half-twist* τ_α *associated to* α is defined by

$$\tau_\alpha = \begin{cases} h \circ \tau \circ h^{-1}(x) & \text{if } x \in N(\alpha) \\ x & \text{if } x \in \Sigma - N(\alpha) \end{cases}$$

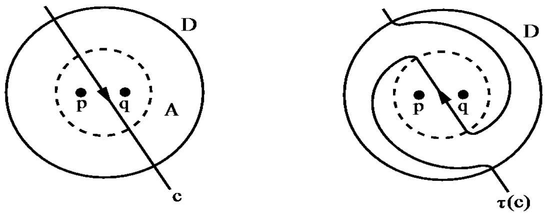

FIGURE 1.8. The basic Dehn half-twist.

Note that the *square* of a Dehn half-twist associated to α is a Dehn twist around a simple closed curve in the interior of a neighborhood of α, encircling once the punctures p and q.

Remark 1.5. In the proof of Theorem 1.4 above, the isomorphism $\mathrm{MCG}(D_n) \to B_n$ is defined by sending the Dehn half-twist associated to the segment joining the j-th and $(j+1)$-th punctures to σ_j.

1.3. **Second geometric definition.** The braid group can be described as a fundamental group of a configuration space. Let M be a manifold. The configuration space of n ordered points in M, or the ordered n-configuration space associated to M, is defined by

$$F_n(M) = \{(x_1, \ldots, x_n) \in M^n \mid x_i \neq x_j \text{ for } i \neq j\} = M^n - \Delta_n,$$

where $\Delta \subset M^n$ is the discriminant, consisting of the union of all hyperplanes $x_i = x_j$ $(i \neq j)$. This configuration space has no manifold structure in the regular sense, but it has a natural stratification structure. The symmetric group S_n acts naturally on $F_n(M)$ by permuting coordinates. The quotient space $C_n(M) = F_n(M)/S_n$ is the configuration space of n *unordered* points in M or simply the unordered n-configuration space associated to M.

Theorem 1.6. *B_n is the fundamental group of $C_n(\mathbb{C})$.*

Proof(sketch). A geometric braid b meets a plane $z = c$ $(0 \leq c \leq 1)$ in exactly n points $(b_n(c), \ldots, b_n(c))$ (first definition). View b as the simultaneous motion of n points in $\mathbb{C}$. Choose a base point $p = (p_1, \ldots, p_n) \in F_n(\mathbb{C})$. Then a braid b defines a loop in $C_n(\mathbb{C})$ with base point $\{b_1(0), \ldots, b_n(0)\} = \{b_1(1), \ldots, b_n(1))\} \in C_n(\mathbb{C})$. It can be shown (by elementary means) that two such loops (with same base point) are homotopic if and only if the associated braids are isotopic. (This was first done by Artin in 1947.) ■

Remark. Fox-Neuwirth (1962) use a cell decomposition of $C_n(\mathbb{C})$ to compute $\pi_1(C_n(\mathbb{C}))$ and show that $C_n(\mathbb{C})$ is a $K(B_n, 1)$- space.

Exercise. Prove that another equivalent definition of B_n is to define B_n as the fundamental group of the space of monic degree n polynomials with complex coefficients and no multiple roots.

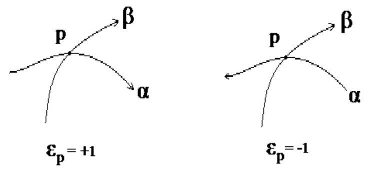

FIGURE 2.1. Local algebraic intersection

2. REPRESENTATIONS

There are many open questions about representations of the braid groups. An intriguing problem is whether braid groups are linear, that is, whether there exists a faithful linear representation of B_n. This problem was solved by D. Krammer first for $n \leq 4$, then by S. Bigelow for all n, who gave a topological proof, and by D. Krammer who gave an algebraic proof.

Our goal in this section is to :

(1) review several basic representations of B_n, from both a topological and an algebraic viewpoint;
(2) prove Krammer–Bigelow's theorem (we plan to present both proofs)
(3) give a survey of known results on and beyond braid groups, and raise some yet unsolved questions.

2.1. **The homological representation.** We start with the simplest non-trivial representation of B_n, the homological representation. We begin by presenting it in topological (homological) terms and then we describe it in purely algebraic terms.

In this paragraph, we view B_n as the mapping class group of the disk D_n with n punctures. Each self-homeomorphism of a connected oriented surface Σ induces an automorphism of $H = H_1(\Sigma; \mathbb{Z})$. We have the commutative diagram

$$\begin{tikzcd} \mathrm{Homeo}(\Sigma) \arrow[r] \arrow[d] & \mathrm{Aut}(H) \\ \mathrm{MCG}(\Sigma) \arrow[ur] & \end{tikzcd}$$

The idea is to look at the action of self-homeomorphisms in homology. Self-homeomorphisms preserve the skew-symmetric bilinear pairing $H \times H \to \mathbb{Z}$ induced by algebraic intersection of (transversal) curves. This intersection pairing is defined as follows. Let $[\alpha], [\beta] \in H$ be represented by smooth simple closed curves. By homotoping β if necessary, we may assume that α and β are transverse. We define

$$[\alpha] \cdot [\beta] = \sum_{p \in \alpha \cap \beta} \varepsilon(x), \quad \text{with } \varepsilon(p) = \begin{cases} +1 & \text{if } T_p\alpha \oplus T_p\beta = T_p\Sigma \\ -1 & \text{if } T_p\alpha \oplus T_p\beta = -T_p\Sigma. \end{cases}$$

The sign is ± 1 according to whether the orientations coincide or not. See Figure 2.1.

Recall that B_n is generated by the Dehn half-twists τ_α. Idea: look at the effect of Dehn twists and Dehn half-twists in homology. First, let α be a simple closed curve in Σ. We compute the action of τ_α in homology as follows. Orient α in an arbitrary way. The effect of τ_α on an oriented curve transversal to α is easily seen from the definition to insert $\pm\alpha$ at each crossing of α with this curve, with ± 1 the sign of the crossing. Namely: for $g \in H$,

$$(\tau_\alpha)_*(g) = g + (\alpha \cdot g)\alpha \tag{2.1}$$

with the abuse of notation consisting in keeping the same notation for the curve α and its homology class in H. Note that the action on H does not depend on the choice of orientation of α.

We now compute the action of τ_α when α is a properly embedded arc, that is, when τ_α is a Dehn half-twist. The figure 1.2 shows the action of τ_α on a curve c. To describe it algebraically, orient the curve α from p to q. We associate to α a loop α' in Σ as follows. Choose a point $z \in \alpha$. Let ℓ_1 (resp. ℓ_2) be the loop in Σ starting at z, moving along α until coming closely q (resp. p), then encircling q (resp. p) in the direction determined by the orientation of Σ and finally moving back to z along α. Set $\alpha' = \ell_1^1\ell_2$. The homotopy class of α' clearly does not depend on the choice of z. The effect of τ_α on an oriented curve transversal to α is to insert $\pm\alpha'$ at each crossing of α with this curve (cf. Figure 1.2). Namely: for $g \in H$,

$$(\tau_\alpha)_*(g) = g + (\alpha \cdot g)\alpha'. \tag{2.2}$$

Here $\alpha \cdot g = -g \cdot \alpha$ is the algebraic intersection number of g with the homology class α of Σ (modulo infinity or with p and q identified).

The viewpoint is fruitful to study the homological representation algebraically, in particular to detect its non-faithfulness. If $\alpha, \beta \subset \Sigma$ are simple closed curves with $\alpha \cdot \beta = 0$, then

$$(\tau_\alpha)_*(\tau_\beta)_*(g) = g + (\alpha \cdot g)\alpha + (\beta \cdot g)\beta = (\tau_\beta)_*(\tau_\alpha)_*(g).$$

This shows that $\alpha \cdot \beta = 0$ implies $[(\tau_\alpha)_*, (\tau_\beta)_*] = 1$ in $\mathrm{Aut}(H)$. On the other hand, we have the following result:

Theorem 2.1. *The Dehn twists τ_α, τ_β commute in $\mathrm{MCG}(\Sigma)$ if and only if α is isotopic to a simple closed curve disjoint from β.*

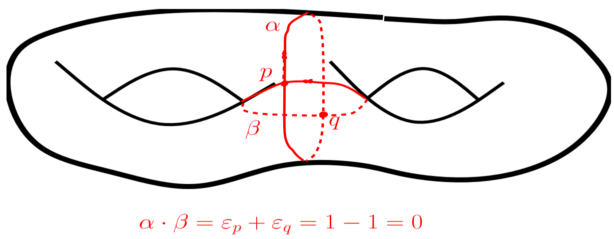

FIGURE 2.2. Not isotopically disjoint simple closed curves with zero algebraic intersection number.

There are simple examples of nontrivial simple closed curves $\alpha, \beta \subset \Sigma$ which are *not* disjoint up to isotopy but have zero algebraic intersection number. See for example Figure 2.1. Then

$$[\tau_\alpha, \tau_\beta] \in \mathrm{Ker}\Big(\mathrm{MCG}(\Sigma, \mathcal{P}) \to \mathrm{Aut}(H)\Big).$$

It follows that the homological representation is not faithful. Similar observations hold for half-twists: if α and β are properly embedded arcs in Σ (i.e., with endpoints in the punctures) such that $\alpha \cdot \beta = 0$, then formula (2.2) implies that $[(\tau_\alpha)_*, (\tau_\beta)_*] = 1$ (easy exercise which follows from the fact that $\alpha \cdot \beta = 0$ is equivalent to $\alpha' \cdot \beta = 0$). Again, there are examples of nontrivial properly embedded arcs $\alpha, \beta \subset \Sigma$ which are not disjoint up to isotopy but have zero algebraic intersection number. See for example Figure 2.1. We obtain similarly as before elements in the kernel of the homological representation.

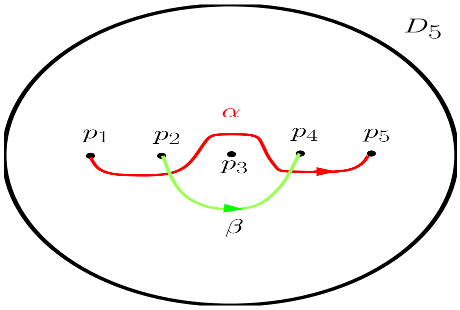

FIGURE 2.3. Not isotopically disjoint properly embedded arcs with zero algebraic intersection number.

Let us consider now the homological representation in the case when the marked surface is D_n, that is, the homological representation of braids $B_n \to \mathrm{Aut}(H) = \mathrm{Aut}(H_1(D_n)) = \mathrm{Aut}(\mathbb{Z}^n)$.

Lemma 2.2. *The homological representation $B_n \to \mathrm{Aut}(\mathbb{Z}^n)$ factorizes through the symmetric group S_n.*

Proof. Fix a basis for $H = \mathbb{Z}^n$ by taking, for each $1 \leq j \leq n$, a small oriented circle e_j around the j-th puncture. We identify σ_j to the Dehn half-twist associated to the segment joining the j-th puncture to the $(j+1)$-th puncture. Using formula (2.2), we see that

$$\sigma_j(e_k) = \begin{cases} e_k & \text{if } j < i \text{ or } j > k+1 \\ e_{k+1} & \text{if } j = k \\ e_{k-1} & \text{if } j+1 = k \end{cases}$$

In terms of matrix, we obtain

$$\sigma_j \mapsto \begin{bmatrix} I_{j-1} & & & 0 \\ & 0 & 1 & \\ & 1 & 0 & \\ 0 & & & I_{n-j+1} \end{bmatrix} = \mathrm{Diag}\Big(I_{j-1}, \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, I_{n-j+1}\Big).$$

This is the standard representation of the symmetric group in terms of permutation matrices. ■

2.2. **The homotopical representation.** We would like to consider finer representations. One natural idea is to use the fundamental group $\pi_1(D_n)$ instead of $H_1(D_n)$. As Artin showed, this is actually the other extreme: we obtain a faithful representation of the braid group as a subgroup of automorphisms of the free group.

Fix a basepoint $z = (0, -1) \in D_n$ once for all. The fundamental group $\pi_1(D_n) = F_n$ the free group of rank n. A system of generators is provided, for example, by a collection of n loops $e_1, \dots, e_n$, starting from z, going up until coming very closely to the j-th puncture ($1 \leq j \leq n$), encircling once in the direction opposite to the one determined by the orientation of D_n, and coming back to z (*). Recall that B_n is the mapping class group of D_n, thus acts on $\pi_1(D_n)$, giving rise thereby to a representation

$$B_n \to \mathrm{Aut}(F_n), \quad \sigma_j \to \widehat{\sigma}_j.$$

To describe more precisely this action, we use the basis $e_1, \dots, e_n$ and compute the action of the half-twists on it. This is done using diagrams of the kind of Figure 1.2. We obtain

$$\widehat{\sigma}_j(e_k) = \begin{cases} e_k & \text{if } k \neq j, j+1 \\ e_j & \text{if } k = j+1 \\ e_j e_{j+1} e_j^{-1} & \text{if } k = j \end{cases}$$

Theorem 2.3 (Artin)**.** *The representation $B_n \to \mathrm{Aut}(F_n)$, $\sigma \mapsto \widehat{\sigma}$ is faithful.*

Artin's original proof is purely combinatorial, see also Birman's book. We will see later a recent nice proof using Dehornoy's ordering. We give here another proof, which takes advantage of the geometric definitions of the braid group.

Proof. The representation given by $\sigma_i \mapsto \widehat{\sigma}_i$ is well-defined, since it is the representation induced by the action of B_n (regarded as the mapping class group) on $\pi_1(D_n)$. In particular, the $\widehat{\sigma}_i$'s verify the braid relations.

To show that this representation is faithful, we shall prove that the image of $\sigma \mapsto \widehat{\sigma}$ is isomorphic to B_n. Denote it, for now, by M_n. We claim that there is a homomorphism $M_n \to B_n$.

By definition, M_n is the subgroup of automorphisms of $\pi_1(D_n)$ which are induced from self-homeomorphisms of D_n which preserve pointwise $\partial D_n = S^1$ and preserve setwise the set of punctures. But any such self-homeomorphism will be isotopic to the identity, by Alexander's trick, and as before, the trace of the isotopy yields a geometric braid (precisely: the image of the punctures times I under the track of the isotopy is a geometric braid). As before, this

(*) Why opposite ? We could as well take the direction determined by the orientation of D_n, but this habit is too well engrained to change it now.

induces a well-defined map $M_n \to B_n$. This map is inverse to the homomorphism $B_n \to M_n \subset \mathrm{Aut}(\pi_1(D_n))$. ■

Remark 2.4. Artin also characterized algebraically all automorphisms φ of F_{n-1} which are in M_n, that is, which come from braids. If $e_1, \dots, e_n$ is a set of generators of F_n, $\beta \in M_n$ if and only if $\beta(e_1 \dots e_n) = e_1 \dots e_n$ and $\beta(e_i)$ is an (arbitrary) conjugate of e_{π_i} where π is a permutation of the indices. Exercise: prove necessity, that is, every automorphism which comes from a braid satisfies the properties stated above.

Therefore, we do have a faithful representation of the braid groups. However, Formanek and Procesi have showed that $\mathrm{Aut}(F_n)$ is *not* linear for $n \geq 3$. Bigelow and Krammer's results can be interpreted as an intermediate representation, between the homological and homotopical representations, which has the best of both worlds: it is linear *and* faithful.

2.3. **The Burau representation.** Before we present Bigelow's solution, we present a long-time candidate for a linear faithful representation of B_n, the Burau representation. As before, we proceed with the topological definition first (which is a generalization of the homological definition) and then describe the Burau representation in purely algebraic terms.

The punctures are denoted as before $p_1, \dots, p_n$ and $z = (0, -1)$ is the base point in D. For each $p \in D$, $H_1(D - p)$ is generated by a small loop encircling p once in the counterclockwise direction and hence is isomorphic to $\mathbb{Z}$. Each loop in $D - p$ represents k times in homology this generator: k is called the *winding number* of the loop around p. There is a unique homomorphism $H_1(D_n) \to \mathbb{Z}$ sending the homology class of a loop to its *total winding number* defined as the sum of its winding numbers around $p_1, \dots, p_n$.

Exercise. Alternatively, work out the definition from the fundamental group $\pi_1(D_n) = F_n$ equipped with its distinguished set of generators.

Let $\widetilde{D}_n \to D_n$ be the normal covering corresponding to the kernel of the total winding number $\phi : \pi_1(D_n) \to \mathbb{Z}$. Note that ϕ is onto. Hence, the group of covering transformations of $\widetilde{D}_n$ is $\mathbb{Z}$, which we write as a multiplicative group with generator t. Set $\Lambda = \mathbb{Z}[t, t^{-1}]$. Our first claim is

Claim 1. The group $H_1(\tilde{D}_n)$ is a Λ-module.

This is just a consequence of the linearization of the action when we pass to homology.

Exercise. $H_1(\tilde{D}_n)$ is a free Λ-module of rank $n - 1$. Hint: this requires a bit of familiarity with the theory of covering spaces and pasting and gluing techniques. Replace punctures by small disks, that is, replace D_n by $X = X_n = D - \cup_i D_i$ where each D_i is a small disk centered at p_i. Clearly the inclusion $X_n \subset D_n$ is a homotopy equivalence. The space X has also a covering $\tilde{X}$ corresponding to the kernel of the total winding number map

ϕ. In particular, it suffices to prove the result for $H_1(\tilde{X})$ (since it will be isomorphic, as a Λ-module, to $H_1(\tilde{D}_n)$). Construct the covering space $\tilde{X}$ of X corresponding to the kernel of the map $\pi_1(X) \to \mathbb{Z}$ (total winding number). There is a description of $\tilde{X}$ from countably many copies of X^* where X_n^* is the space obtained from X_n by removing a segment I_i from p_i to the boundary. We hence have two copies I_i^+ and I_i^- of I in X_n^*. The space $\tilde{X}$ is obtained by taking the disjoint union $\coprod_{j\in\mathbb{Z}} X_j^*$ of countably many copies of X^* and identifying $I_i^+ \subset X_j^*$ to $I_i^- \subset X_{j+1}^*$. The action of $t \in \mathbb{Z} = \mathrm{Aut}(\tilde{X})$ can be thought of as moving one point in $\tilde{X}$ one level upward. Clearly the orbit space is X. Let A (resp. B) the subspace of $\tilde{X}$ induced by $\cup_{j\in\mathbb{Z}} X_{2j}^*$ (resp. by $\cup_{j\in\mathbb{Z}} X_{2j+1}^*$). The homology of $H_1(\tilde{X})$ can be computed using a Mayer–Vietoris argument, applied to A and B. A basis of $H_1(\tilde{X})$ consists of $v_1, \ldots, v_{n-1}$ where v_i is induced by the lift of the loop $e_i e_{i+1}^{-1}$. ∎

Choose a lift $\tilde{z} \in \tilde{D}_n$ of the base point $z \in \partial D_n$.

Theorem 2.5. *There is a well-defined representation* $B_n \to \mathrm{Aut}(H_1(\tilde{D}_n))$,

Proof. Let $\beta = h_* \in B_n$ be induced by a self-homeomorphism h of D_n. For any loop $\gamma \in D_n$ based at z, we have $\phi(h \circ \gamma) = \phi(\gamma)$ (Why ?). It follows that h lifts to a self-homeomorphism of $\tilde{D}_n$, thereby gives rise to an automorphism $\tilde{h}_*$ of $H_1(\tilde{D}_n)$. If h' is another self-homeomorphism such that $h'_* = \beta$, then since $\pi_1(\tilde{D}_n) \to \pi_1(D_n)$ is injective, we have $\tilde{h}'_* = \tilde{h}_*$ as automorphisms of $\pi_1(\tilde{D}_n)$. In particular, they coincide as automorphisms of $H_1(\tilde{D}_n)$. ∎

Exercise. Prove carefully the statement $\phi(h \circ \gamma) = \phi(\gamma)$ in the proof above. Hint: it amounts to proving that self-homeomorphisms of D_n (applied to loops) preserve the total winding number 0. Compare with Remark 2.4.

W. Burau introduced the representation $\beta = h_* \mapsto \tilde{h}_*$ which bears his name in 1935.

Exercise. Describe the Burau representation in terms of matrices (for this, express the action of Dehn half-twists, on the basis of $H_1(\tilde{D}_n)$ given above) and show that it is irreducible. For this, describe the action of the Artin generators on $H_1(\tilde{D}_n)$. Answer:

$$(\tilde{\sigma}_i)_*(v_j) = \begin{cases} v_j + t v_{j+1} & \text{if } j = i-1 \\ -t v_j & \text{if } j = i \\ v_{j-1} + v_j & \text{if } j = i+1 \\ v_j & \text{otherwise} \end{cases}$$

The classical Burau representation is given by

$$\sigma_i \mapsto \begin{bmatrix} I_{j-1} & & & 0 \\ & 1-t & t & \\ & 1 & 0 & \\ 0 & & & I_{n-j+1} \end{bmatrix} = \mathrm{Diag}\Big(I_{j-1}, \begin{bmatrix} 1-t & t \\ 1 & 0 \end{bmatrix}, I_{n-j+1}\Big).$$

Exercise. Show that the classical Burau representation splits as the Burau representation given above and a trivial representation of dimension 1.

Exercise. Show that all nontrivial representations of B_n of the form $\sigma_i \mapsto \mathrm{Diag}\Big(I_{i-1}, \begin{bmatrix} a & b \\ c & d \end{bmatrix}, I_{n-i+1}\Big)$ are determined by the Burau representation.

The ideas above show that the Burau representation is a refinement of the homological representation. So it may not be surprising that the algebraic intersection pairing (defined in §2.1) can also be refined as an intersection pairing on $\tilde{D}_n$ with values in Λ.

Let α, β be two properly embedded oriented arcs in D_n (with endpoints in the punctures). Lift α, β to $\tilde{\alpha}$ and $\tilde{\beta}$ in $\tilde{D}_n$. We define

$$\langle \alpha, \beta \rangle = \sum_{k \in \mathbb{Z}} (t^k \tilde{\alpha} \cdot \tilde{\beta}) t^k \in \Lambda/(t). \tag{2.3}$$

Here $t^k\tilde{\alpha}$ denotes the image of $\tilde{\alpha}$ in $\tilde{D}_n$ by the covering automorphism t^k. The algebraic intersection number of $t^k\tilde{\alpha}$ and $\tilde{\beta}$ in $\tilde{D}_n$ is denoted $t^k\tilde{\alpha} \cdot \tilde{\beta}$ and is an integer.

Remark 2.6. The sum is finite. [Proof: we claim that $t^k\tilde{\alpha} \cdot \tilde{\beta} = 0$ for k large enough. Otherwise, since $\tilde{\beta}$ is compact, the set $\{t^k\tilde{\alpha} \cap \tilde{\beta}\}_k$ would have a limit point in $\tilde{\beta}$. This is ruled out by transversality.]

Remark 2.7. Dependence on the choice of the lift. Since the covering is normal, the automorphism group $\mathbb{Z}$ acts transitively on the fibre. Hence the lifts are determined up to a translation by t^k. Replacing $\tilde{\alpha}$ by $t^k\tilde{\alpha}$ and $\tilde{\beta}$ by $t^l\tilde{\beta}$ and using the fact that $t\tilde{\alpha} \cdot \tilde{\beta} = \tilde{\alpha} \cdot t^{-1}\tilde{\beta}$, we see that $\langle \alpha, \beta \rangle$ is well-defined up to multiplication by a power of t.

Remark 2.8. Denote by bar the involution of Λ sending t to t^{-1}. We have $\langle \alpha, \beta \rangle = -\overline{\langle \beta, \alpha \rangle}$.

Exercise. A practical computation of the refined intersection pairing. Show that

$$\langle \alpha, \beta \rangle = \sum_{p \in \alpha \cap \beta} \varepsilon_p t^{k_p},$$

where ε_p is the intersection sign at $p \in \alpha \cap \beta$ and $k_p \in \mathbb{Z}$ is defined by

$$k_p - k_q = \phi(c^{\alpha\beta}_{pq}),$$

where $c^{\alpha\beta}_{pq}$ is the loop starting at p going to q along α and then from q to p along β.

As in the case of the homological representation, one can show

Proposition 2.9. *If $\langle \alpha, \beta \rangle = 0$, then $(\tilde{\tau}_\alpha)_*$ and $(\tilde{\tau}_\beta)_*$ commute in $H_1(\tilde{D}_n)$.*

Using this result, S. Bigelow showed that B_5 is not faithfully represented by the Burau representation. This implies that B_n is not faithfully represented for $n \geq 5$ (this was proved earlier for $n \geq 9$, using different methods by J. Moody). The Burau representation is known to be faithful for $n = 3$. Whether the Burau representation of B_n is faithful for $n = 4$ remains open.

2.4. **Lawrence–Krammer representation.** It should be by now clear that we need to find a more sophisticated covering for D_n, if we want a faithful linear representation.

We continue using the notation of D, P, D_n introduced in the previous sections. Let $C = C_k(D_n)$ be the space of unordered k-tuples of distinct points in D_n. This space can be formally defined as $(D_n^k - \Delta_k)/S_k$, where Δ_k is the subset of k-tuples with at least two equal coordinates and S_k is the symmetric group permuting the coordinates. The space C is a connected noncompact 2^k-manifold. The orientation for C is induced by the orientation of D_n. A loop α in C is a map $\alpha : I \to C$ such that $\alpha(0) = \{\alpha_1(0), \dots, \alpha_k(0)\} = \{\alpha_1(1), \dots, \alpha_k(1)\}$. Each α_j is a arc or a loop in C.

Lemma 2.10. *The image in D_n of*

$$M = \bigcup_{\substack{s\in I \\ 1\leq j\leq k}} \alpha_j(s)$$

is a disjoint union of at most k oriented circles.

Denote by $a_\alpha \in \mathbb{Z}$ the total winding number of this 1-manifold M around the punctures $p_1, \dots, p_n \in P$. For $i < j$, consider the map $\tilde{l}_{ij}$ defined

$$I = [0,1] \to S^1, \quad s \mapsto \frac{\alpha_i(s) - \alpha_j(s)}{|\alpha_i(s) - \alpha_j(s)|}.$$

Suppose that $\{\alpha_i, \alpha_j\}$ is known: the map $\tilde{l}_{ij}$ is well-defined up to sign. Thus the composition $l_{ij} = p \circ \tilde{l}_{ij}$ where $p : S^1 \to S^1/\{\pm 1\} = \mathbb{R}P^1 = S^1$ is the canonical (antipodal) projection, is well-defined. However, we may not have $\alpha_i(0) - \alpha_j(0) = \alpha_i(1) - \alpha_j(1)$. We do know that $\alpha_i(1) - \alpha_j(1) = \alpha_{\pi(i)}(0) - \alpha_{\pi(j)}(0)$ for some permutation $\pi \in S_k$. In other words, we need to take into account the induced action of S_k on the $\binom{k}{2}$ tuples $\alpha_i(0) - \alpha_j(0)$, $1 \leq i < j \leq k$. Denote by T the space $C_{\binom{k}{2}}(S^1)$ of $\binom{k}{2}$ unordered tuples of distinct points in S^1.

The idea is to consider the map l

$$I = [0,1] \to T, \quad s \mapsto \Big\{ l_{ij}(s) \Big\}_{1\leq i<j\leq k}.$$

Observe that l is actually a loop in T. Therefore, there is an induced element $b_\alpha = [l] \in H_1(T; \mathbb{Z})$.

Lemma 2.11. $H_1(T; \mathbb{Z}) = H_1(S^1; \mathbb{Z}) = \mathbb{Z}$.

Define a homomorphism $\phi : H_1(C) \to \mathbb{Z}\langle q, t\rangle$ by sending $\alpha \to q^{a_\alpha} t^{b_\alpha}$. Here $\mathbb{Z}\langle q, t\rangle$ denotes the free abelian group on two (multiplicative) generators q, t. Denote by Λ the group ring $\mathbb{Z}[q^{\pm 1}, t^{\pm 1}]$. Consider the regular covering $\tilde{C} \to C$ associated to Ker ϕ. By the general theory of covering spaces, $\mathrm{Aut}(\tilde{C})$ is $\mathbb{Z}^2$, generated by q and t. It follows that $H_2(\tilde{C})$ becomes a Λ-module.

Lemma 2.12. *Any self-homeomorphism $h : D_n \to D_n$ induces a self-homeomorphism $\tilde{h} : \tilde{C} \to \tilde{C}$ commuting with the covering transformations.*

Proof. Clearly h induces a self-homeomorphism $h : C \to C$ by $h(\{x, y\}) = \{h(x), h(y)\}$. Set c be a basepoint in C such that $c_j \in \partial D$ for all j. Then $h(c) = c$ (since h is the identity on the boundary). Furthermore, h acts on $H_1(C)$ and the action commutes with ϕ (exercise). It follows tha $h : C \to C$ lifts uniquely to a map $\tilde{h} : \tilde{C} \to \tilde{C}$ which fixes the fibre over c pointwise and commutes with the covering transformations. ■

Consequently, there is an induced Λ-module automorphism $\tilde{h}_*$ on homology $H_*(\tilde{C})$. From now on, we assume that $\mathbf{k} = \mathbf{2}$. The Lawrence–Krammer representation is the map

$$\rho : B_n \to \mathrm{Aut}(H_2(\tilde{C}), \quad \sigma = [h] \mapsto \tilde{h}_*.$$

3. Bigelow's work

We present Bigelow's proof of the faithfulness of the Lawrence–Krammer representation. We must show that the kernel of ρ is trivial. Bigelow's method consists in defining a certain pairing between forks and noodles that will preserved by any element in the kernel.

3.1. **Forks and noodles.** The terminology is due to D. Krammer. A *fork* is an embedded tree $F \subset D$, with four vertices and three edges: the vertices are the basepoint $d \in \partial D$, two punctures p_i, p_j and a point $z \in D - P$ which is the common vertex to the three edges; the only intersection between the edges and ∂D is d and the only intersection between the edges and $P = \{p_1, \ldots, p_n\}$ is $\{p_i, p_j\}$. The *handle* $H(F)$ of the fork F is the edge joining d to z. The *tin* $T(F)$ of the fork F is the union of $F - H(F)$ and z.

Using the orientation of D, it is always possible to create a parallel copy of a fork, by pushing slightly the fork F to its distinguished side (fixing the punctures and moving slightly the basepoint on the boundary in the counterclockwise direction).

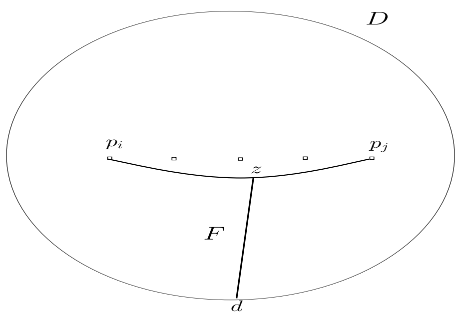

FIGURE 3.1. A fork.

To each fork $F \subset D$ is associated a surface $\tilde{\Sigma}(F) \subset \tilde{C}$ as follows. Let F' be a parallel copy of F, with basepoint d' (close to d) and common edge z' (close to z). Let $\Sigma(F)$ be pairs $\{y, y'\} \in C$ such that $y \in T(F) - \{p_i, p_j\}$ and $y' \in T(F') - \{p_i, p_j\}$. Note that $\Sigma(F)$ is homeomorphic to the product of two copies of the open unit interval. We now use the handle to lift $\Sigma(F)$ in $\tilde{C}$ as follows. Fix once for all a basepoint $\tilde{c_0} \in \tilde{C}$ over $c_0 = \{z, z'\} \in C$. Let β_1 be the arc from d to z along $H(F)$. Similarly let β_1' be the arc from d' to z' along $H(F')$. There is a unique lift $\tilde{\beta}$ of $\{\beta_1, \beta_1'\}$ to $\tilde{C}$ which starts at $\tilde{c_0}$. Let $\tilde{\Sigma}(F)$ be the lift of $\Sigma(F)$ to $\tilde{C}$ that contains the lift $\tilde{\beta}(1)$ of $\{z, z'\}$. The orientations of $\Sigma(F)$ and $\tilde{\Sigma}(F)$ are determined by the orientation of $T(F), T(F')$.

A *noodle* is an embedded edge $N \subset D - P$ with endpoints d, d', oriented from d to d'.

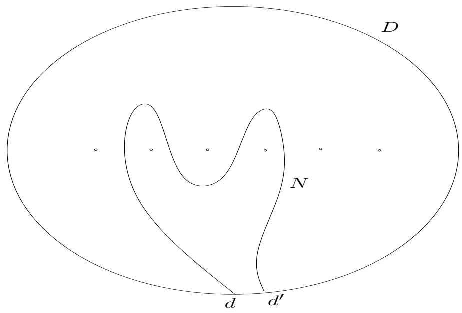

FIGURE 3.2. A noodle.

Let $\Sigma(N)$ be the set of pairs $\{y, y'\} \in C$ such that $y \neq y'$ and $y, y' \in N$. This is a surface which can be oriented: for $\{x, y\} = \{y, x\} \in \Sigma(N)$, either x or y is closer to d along N. Say x is closer; then for the orientation of $\Sigma(N)$, take the product orientation of N at x and y in this order.

Finally, we take $\tilde{\Sigma}(N)$ to be the lift of $\Sigma(N)$ to $\tilde{C}$ by requiring it to contain $\tilde{c_0} \in \tilde{C}$. Orientation lifts as usual and we have $\tilde{\Sigma}(N) \cap \partial\tilde{C} = \partial\tilde{\Sigma}(N)$.

3.2. **Product between forks and noodles.** Let F be a fork, N a noodle. We define

$$\langle N, F\rangle = \sum_{(a,b)\in\mathbb{Z}^2} q^a t^b (q^a t^b \tilde{\Sigma}(N), \tilde{\Sigma}(F)) \in \Lambda = \mathbb{Z}[q^{\pm 1}, t^{\pm 1}].$$

Here $q^a t^b \tilde{\Sigma}(N)$ is the image of $\tilde{\Sigma}(N)$ under the covering transformation $q^a t^b$; the symbol $(\cdot, \cdot)$ denotes the algebraic intersection number (with values in $\mathbb{Z}$) in $\tilde{C}$. Several questions arise:

(1) Is it finite ? Is it well-defined ?
(2) Is it invariant under isotopies of N and F in D_n (fixed on the endpoints) ?

S. Bigelow proves that the answer to all these questions is positive. The main problem is that an algebraic intersection number between two surfaces, even properly embedded and in general position, may not be well-defined since it may be possible to eliminate intersecions by pushing them off to infinity. Bigelow circumvents this problem by exhibiting a *closed* immersed surface representing a certain multiple of the original outside a small neighborhood of the punctures (the algebraic intersection number between a properly embedded surface and an immersed closed surface is well-defined). This is done by a direct computation and we refer to the original article for details. In fact, Krammer gives another method for computing $H_2(\tilde{C})$.

The following lemma is a direct consequence of the definitions, once the pairing is seen to be well-defined.

Lemma 3.1. *Let h be a self-homeomorphism representing an element in the kernel of the Lawrence–Krammer representation. Then $\langle N, h(F)\rangle = \langle N, F\rangle$ for any fork F and any noodle N.*

The rest of Bigelow's proof relies on two fundamental lemmas, the first of which is his own hard work. The second is a classical tool of the topology of surfaces.

Lemma 3.2 (Key Lemma)**.** *Let N be a noodle and F a fork. Then $\langle N, F\rangle = 0$ if and only if $T(F)$ is isotopic relative to $\partial D \cup P$ to an arc disjoint from N.*

Of course one direction (only if) is easy. The hard part consists in proving that the algebraic condition on the pairing determines the geometric conclusion.

Lemma 3.3 (Basic Lemma)**.** *Let α, β be simple closed curves in D_n that intersect transversely at finitely many points $P = \{p_1, \ldots, p_n\}$. The following assertions are equivalent:*

1. *α is isotopic relative to $\partial D \cup P$ to a simple closed curve that intersects β at fewer points.*
2. *α and β cobound a bigon (= embedded disc in D_n whose boundary consists of one subarc of α and one subarc of β).*

3.3. **A proof of the basic lemma.** One direction is easy: if the curves cobound a bigon, then α can be isotoped so as to eliminate two intersection points with β. See Figure 3.3.

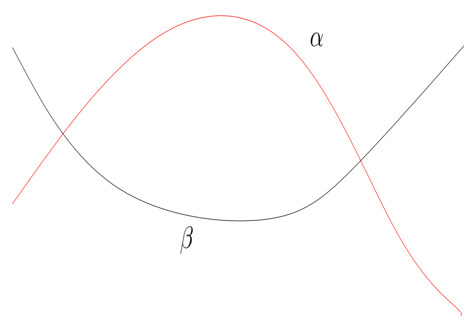

FIGURE 3.3. A bigon between two curves.

The converse is more interesting. We must show that if the cardinality of $\alpha \cap \beta$ is not minimal among simple closed curves isotopic to α and β (exercise: define this precisely), then there exists two points p_1, p_2 of $\alpha \cap \beta$

and two paths γ_0, γ_1 joining p_1 to p_2 respectively along α and β such that the (singular) loop $\gamma_0 * \gamma_1^{-1}$ is homotopically trivial in D_n.

By hypothesis, there exists an isotopy $h_t : S^1 \to D - P$ such that h_0 parametrizes α and such that $h_1(S^1) \cap \beta$ has fewer points than $\alpha \cap \beta$. We may assume the isotopy transverse to β, that is, $h : S^1 \times [0,1] \to D - P$ is transverse to β. Then $h^{-1}(\beta)$ is a submanifold of dimension 1, transverse to the boundary. The intersection points $p_1, \ldots, p_n$ are the elements of $h_0^{-1}(\alpha \cap \beta)$. There are at most four types (I, II, III and IV) of connected components of $h^{-1}(\beta)$. By hypothesis, there is at least one component γ_1 of type I (since the number of intersection points must decrease during the isotopy). Let γ_0 be the arc joining p_1 to p_2 along $S^1 \times 0$. Once the correct orientations are chosen, we clearly have $\gamma_0 * \gamma_1^{-1}$ homotopically trivial in $D - P$. ∎

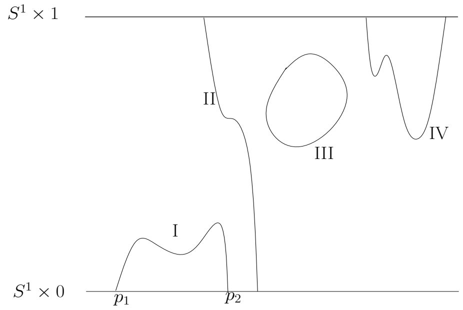

FIGURE 3.4. A schematic view of $h^{-1}(\beta)$.

3.4. **An idea of the proof of the key lemma.**

3.5. **End of proof.** We deduce the theorem (faithfulness of the Lawrence representation) from the basic lemma and the key lemma. Let h be a self-homeomorphism of $D - P$ representing an element in the kernel of $\rho : B_n \to \mathrm{Aut}(H_2(\tilde{C})$. We must show that h is isotopic to the identity through self-homeomorphisms preserving pointwise the boundary and setwise the punctures.

We begin with a direct consequence of the key lemma: An embedded arc T in $D - P$ with endpoints in distinct punctures can be isotoped off a noodle N if and only if $h(T)$ can be isotoped off N. Indeed, extend T to a fork F (so that T is the tine of F). Lemma 3.1 says that $\langle N, h(F) \rangle = 0$ if and only if $\langle N, F \rangle = 0$. Then the key lemma applies.

We now apply this observation to the following arcs and noodles. Let T_j be the segment joining p_j to p_{j+1}, $j = 1, \ldots, n-1$. Denote by N_j the noodle starting from d towards p_j, encircling p_j in the clockwise direction and then returning immediately to d'. Clearly T_j can be isotoped off the noodle N_k if and only if $k \neq j, j+1$. Hence, h must induce the identity on the set P of punctures.

Since T_1 is disjoint from N_3, there is an isotopy of h ensuring that $h(T_1)$ is disjoint from N_3. This does not change the class $[h]$ in the kernel of the representation, so we keep the same letter h for the new self-homeomorphism. Similarly, $h(T_1)$ can be made disjoint from N_4. Now apply the basic lemma to ensure that this is done by a sequence of isotopies eliminating bigons for the pair $(N_4, h(T_1))$. Since N_4 and $h(T_1)$ do not meet N_3, neither do these bigons; thus there are no new intersections created, during the isotopy, between N_3 and $h(T_1)$. Repeat this argument until $h(T_1)$ is made disjoint from all N_i, $i = 3, 4, \ldots, n-1$. Apply another isotopy so that $h(T_1) = T_1$.

Now apply the same procedure to T_2 so that $h(T_2) = T_2$ while keeping $h(T_1) = T_1$. Continue until $h(T_i) = T_i$ for all $i = 1, \ldots, n-1$. It is not hard to see that such a self-homeomorphism is isotopic to a power of a Dehn twist around a circle parallel to ∂D enclosing all the punctures (exercise, using the fact that elements of the mapping class group are determined by their diagrams on D_n). Such a Dehn twist represents the Garside braid Δ^2 where Δ is inductively defined by $\Delta_2 = \sigma_1$, $\Delta_3 = \sigma_1\sigma_2\sigma_1$ and (n being the number of strands)

$$\Delta_n = \sigma_1\sigma_2 \ldots \sigma_n - 1 \cdot \Delta_{n-1}.$$

Such a twist acts on $H_2(\tilde{C})$ by multiplication by $q^{2n}t^2$. Hence h acts as $(q^{2n}t^2)^k$ for some $k \in \mathbb{Z}$. But by hypothesis, h acts as the identity. Therefore $k = 0$ and h represents the trivial braid. ■

Exercise. Prove the assertion about the action of Δ^2 on $H_2(\tilde{C})$, assuming that all elements in $H_2(\tilde{C})$ can be represented by a noodle.

3.6. **Remark.** In a recent paper, S. Bigelow defines a purely homological pairing that somehow bypasses the indirect argument he gives to show that the pairing between noodles and forks is well-defined.

Florian Deloup, Laboratoire Emile Picard, UMR 5580 CNRS/Université Paul Sabatier, 118 route de Narbonne, 31062 Toulouse, France.

E-mail address: deloup@picard.ups-tlse.fr