Braid Groups and Iwahori-Hecke Algebras

Stephen Bigelow

ABSTRACT. The braid group B_n is the mapping class group of an n-times punctured disk. The Iwahori-Hecke algebra $\mathcal{H}_n$ is a quotient of the braid group algebra of B_n by a quadratic relation in the standard generators. We discuss how to use $\mathcal{H}_n$ to define the Jones polynomial of a knot or link. We also summarize the classification of the irreducible representations of $\mathcal{H}_n$. We conclude with some directions for future research that would apply mapping class group techniques to questions related to $\mathcal{H}_n$.

1. Introduction

The braid group B_n is the mapping class group of an n-times punctured disk. It can also be defined using certain kinds of arrangements of strings in space, or certain kinds of diagrams in the plane. Our main interest in the braid group B_n will be in relation to the Iwahori-Hecke algebra $\mathcal{H}_n$, which is a certain quotient of the group algebra of B_n. The exact definition will be given in Section 3.

The Iwahori-Hecke algebra plays an important role in representation theory. It first came to the widespread attention of topologists when Jones used it to define the knot invariant now called the Jones polynomial [Jon85]. This came as a huge surprise, since it brought together two subjects that were previously unrelated. It has given knot theorists a host of new knot invariants, and intriguing connections to other areas of mathematics to explore. It has also helped to promote the use of pictures and topological thinking in representation theory.

As far as I know, no major results related to the Iwahori-Hecke algebra have yet been proved using the fact that B_n is a mapping class group. I think the time is ripe for such a result. Unfortunately this paper will use the diagrammatic definition of the braid group almost exclusively. I hope it will at least help to provide a basic grounding for someone who wants to pursue the connection to mapping class groups in the future.

The outline of this paper is as follows. In Sections 2 and 3 we introduce the braid group and Iwahori-Hecke algebra. In Sections 4 and 5 we give a basis for the Iwahori-Hecke algebra and for its module of trace functions. In Section 6 we explain how one such trace function leads to the definition of the Jones polynomial of a knot or a link. In Section 7 we briefly summarize the work of Dipper and James [DJ86]
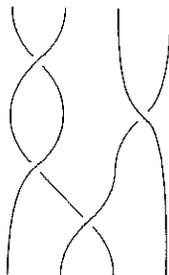
FIGURE 1. A braid with four strands.

classifying the irreducible representations of the Iwahori-Hecke algebra. In Section 8 we conclude with some speculation on possible directions for future research. Open problems will be scattered throughout the paper.

2. The braid group

Like most important mathematical objects, the braid group B_n has several equivalent definitions. Of greatest relevance to this volume is its definition as a mapping class group. Let D be a closed disk, let $p_1, \ldots, p_n$ be distinct points in the interior of D, and let $D_n = D \setminus \{p_1, \ldots, p_n\}$. The braid group B_n is the mapping class group of D_n. Thus a braid is the equivalence class of a homeomorphism from D_n to itself that acts as the identity of the boundary of the disk.

Artin's original definition of B_n was in terms of *geometric braids*. A geometric braid is a disjoint union of n edges, called *strands*, in $D \times I$, where I is the interval $[0, 1]$. The set of endpoints of the strands is required to be $\{p_1, \ldots, p_n\} \times \{0, 1\}$, and each strand is required to intersect each disk cross-section exactly once. Two geometric braids are said to be equivalent if it is possible to deform one to the other through a continuous family of geometric braids. The elements of B_n are equivalence classes of geometric braids.

We will need some terminology to refer to directions in a geometric braid. Take D to be the unit disk centered at 0 in the complex plane. Take $p_1, \ldots, p_n$ to be real numbers with $-1 < p_1 < \cdots < p_n < 1$. The *top* and *bottom* of the braid are $D \times \{1\}$ and $D \times \{0\}$, respectively. In a disk cross-section, the *left* and *right* are the directions of decreasing and increasing real part, respectively, while the *front* and *back* are the directions of decreasing and increasing imaginary part, respectively.

Multiplication in B_n is defined as follows. If a and b are geometric braids with n strands then the product ab is obtained by stacking a on top of b and then rescaling vertically to the correct height. This can be shown to give a well-defined product of equivalence classes, and to satisfy the axioms of a group.

A geometric braid can be drawn in the plane using a projection from $D \times I$ to $[-1, 1] \times I$. An example is shown in Figure 1. The projection map is given by $(x + iy, t) \mapsto (x, t)$. Note that this sends each strand to an embedded edge. We also require that the braid be in *general position* in the sense that the images of the strands intersect each other transversely, with only two edges meeting at each point of intersection. The points of intersection are called *crossings*. At each crossing, we record which of the two strands passed in front of the other at the corresponding disk cross-section of the geometric braid. This is usually represented pictorially by
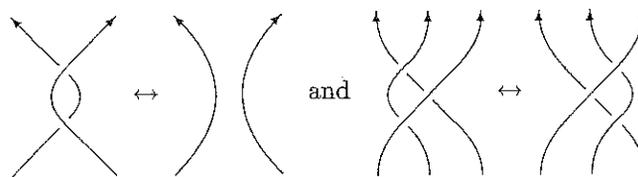
FIGURE 2. Reidemeister moves of types two and three.

a small break in the segment that passes behind. The image of a geometric braid under a projection in general position, together with this crossing information, is called a *braid diagram*.

Let us fix some terminology related to braid diagrams. The directions left, right, top, and bottom are the images of these same directions in the geometric braid, so for example the point $(-1, 1)$ is the top left of the braid diagram. We say a strand makes an *overcrossing* or an *undercrossing* when it passes respectively in front of or behind another strand at a crossing. A crossing is called *positive* if the strand making the overcrossing goes from the bottom left to the top right of the crossing, otherwise it is called *negative.* The endpoints of strands are called *nodes.*

Two braid diagrams represent the same braid if and only if they are related by an isotopy of the plane and a sequence of Reidemeister moves of types two and three. These are moves in which the diagram remains unchanged except in a small disk, where it changes as shown in Figure 2. (There is also a Reidemeister move of type one, which is relevant to knots but not to braids.)

For $i = 1, \ldots, n-1$, let σ_i be the braid diagram with one crossing, which is a positive crossing between strands i and $i+1$. The braid group B_n is generated by $\sigma_1, \ldots, \sigma_{n-1}$, with defining relations

- $\sigma_i \sigma_j = \sigma_j \sigma_i$ if $|i - j| > 1$,
- $\sigma_i \sigma_j \sigma_i = \sigma_j \sigma_i \sigma_j$ if $|i - j| = 1$.

There is an imprecise but vivid physical description of the correspondence between a geometric braid and a mapping class of the n-times punctured disk D_n. Imagine a braid made of inflexible wires and a disk made of flexible rubber. Press the disk onto the top of the braid, puncturing the disk at n points in its interior. Now hold the disk by its boundary and push it down. As the wires of the braid twist around each other, the punctures of the disk will twist around and the rubber will be stretched and distorted to accommodate this. The mapping class corresponding to the geometric braid is represented by the function taking each point on D_n to its image in D_n after the disk has been pushed all the way to the bottom of the braid. (With our conventions, this description gives the group of mapping classes acting on the right.)

See [**Bir74**], or [**BB**], for proofs that these and other definitions of B_n are all equivalent. This paper will primarily use the definition of a braid group as a braid diagram. This is in some sense the least elegant choice since it involves an arbitrary projection and a loss of the true three-dimensional character of the geometric braid. The main goal is to provide an introduction that may inspire someone to apply mapping class group techniques to problems that have previously been studied algebraically and combinatorially.
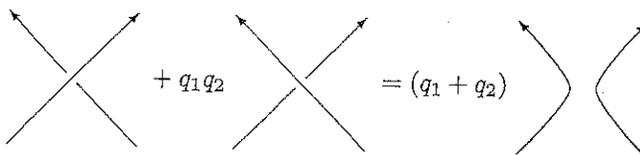
FIGURE 3. The skein relation

3. The Iwahori-Hecke algebra

Let n be a positive integer and let q_1 and q_2 be units in a domain R. The *Iwahori-Hecke algebra* $\mathcal{H}_n(q_1, q_2)$, or simply $\mathcal{H}_n$, is the associative R-algebra given by generators $T_1, \ldots, T_{n-1}$ and relations

- $T_i T_j = T_j T_i$ if $|i - j| > 1$,
- $T_i T_j T_i = T_j T_i T_j$ if $|i - j| = 1$,
- $(T_i - q_1)(T_i - q_2) = 0$.

The usual definition of the Iwahori-Hecke algebra uses only one parameter q. It corresponds to $\mathcal{H}_n(-1, q)$, or in some texts to $\mathcal{H}_n(1, -q)$. There is no loss of generality because there is an isomorphism from $\mathcal{H}_n(q_1, q_2)$ to $\mathcal{H}_n(-1, -q_2/q_1)$ given by $T_i \mapsto -q_1 T_i$. It will be convenient for us to keep two parameters.

We now explore some of the basic properties of the Iwahori-Hecke algebra. Since q_1 and q_2 are units of R, the generators T_i are units of $\mathcal{H}_n$, with

$$T_i^{-1} = (T_i - q_1 - q_2)/(q_1 q_2).$$

Thus there is a well-defined homomorphism from B_n to the group of units in $\mathcal{H}_n$ given by

$$\sigma_i \mapsto T_i.$$

The following is a major open question.

QUESTION 1. *If $R = \mathbf{Q}(q_1, q_2)$, is the above map from B_n to $\mathcal{H}_n(q_1, q_2)$ injective?*

For $n = 3$, the answer is yes. For $n = 4$, the answer is yes if and only if the Burau representation of B_4 is injective, or *faithful*. The Burau representation is one of the irreducible summands of the Iwahori-Hecke algebra over $\mathbf{Q}(q_1, q_2)$. By a result of Long [**Lon86**], the map from B_n to $\mathcal{H}_n$ is injective if and only if at least one of these irreducible summands is faithful. For $n = 4$, they are all easily shown to be unfaithful except for the Burau representation, which remains unknown. For $n \geq 5$, the Burau representation is unfaithful [**Big99**], but there are other summands whose status remains unknown.

One can also ask Question 1 for other choices of ring R and parameters q_1 and q_2. If the map from B_n to $\mathcal{H}_n$ is injective for any such choice then it is injective when $R = \mathbf{Q}(q_1, q_2)$. A non-trivial case when the map is not injective is when $n = 4$ and $R = k[q_1^{\pm 1}, q_2^{\pm 1}]$, where k is a field of characteristic 2 [**CL97**] or 3 [**CL98**]. Another is when $n = 4$, $R = \mathbf{Q}$ and $q_2/q_1 = -2$ [**Big02**].

Using the map from B_n to $\mathcal{H}_n$, we can represent any element of $\mathcal{H}_n$ by a linear combination of braid diagrams. The quadratic relation is equivalent to the *skein relation* shown in Figure 3. Here, an instance of the skein relation is a relation involving three diagrams that are identical except inside a small disk where they are as shown in the figure.

One motivation for studying the Iwahori-Hecke algebra is its connection with the representation theory of the braid groups. The representations of $\mathcal{H}_n$ are precisely those representations of B_n for which the image of the generators satisfy a quadratic relation. The study of these representations led Jones to the discovery of his knot invariant, which we define in Section 6

Another reason for interest is the connection between the Iwahori-Hecke algebra and the symmetric group. There is an isomorphism from $\mathcal{H}_n(1,-1)$ to the group algebra $R\mathfrak{S}_n$ taking T_i to the transposition $(i, i+1)$. Thus $\mathcal{H}_n(q_1, q_2)$ can be thought of as a *deformation* of $R\mathfrak{S}_n$. The Iwahori-Hecke algebra plays a role in the representation theory of the general linear group over a finite field that is analogous to the role of the symmetric group in the representation theory of the general linear group over the real numbers. See for example [Dip85].

This process of realizing a classical algebraic object as the case $q = 1$ in a family of algebraic objects parametrized by q is part of a large circle of ideas called *quantum mathematics*, or *q-mathematics*. The exact nature and significance of any connection to quantum mechanics not clear at present. One example is [Bar03], in which Barrett uses quantum mathematics to analyze quantum gravity in a universe with no matter and three space-time dimensions.

4. A basis

The aim of this section is to show that $\mathcal{H}_n$ is a free R-module of rank $n!$, and to give an explicit basis.

Let $\phi: B_n \to \mathfrak{S}_n$ be the map such that $\phi(\sigma_i)$ is the transposition $(i, i+1)$. Thus in any braid b, the strand with lower endpoint at node number i has upper endpoint at node number $\phi(b)(i)$.

For $w \in \mathfrak{S}_n$, let T_w be a braid diagram with the minimal number of crossings such that every crossing is positive and $\phi(T_w) = w$. Such a braid can be thought of as "layered" in the following sense. In the front layer is a strand connecting node 1 at the bottom to node $w(1)$ at the top. Behind that is a strand connecting node 2 at the bottom to node $w(2)$ at the top. This continues until the back layer, in which a strand connects node n at the bottom to node $w(n)$ at the top. From this description it is clear that our definition of T_w specifies a unique braid in B_n. By abuse of notation, let T_w denote the image of this braid in $\mathcal{H}_n$. For example, if w is a transposition $(i, i+1)$ then T_w is the generator T_i.

THEOREM 4.1. *The set of T_w for $w \in \mathfrak{S}_n$ forms a basis for $\mathcal{H}_n$.*

To prove this, we first describe an algorithm that will input a linear combination of braid diagrams and output a linear combination of basis elements T_w that represents the same element of $\mathcal{H}_n$. By linearity, it suffices to describe how to apply the algorithm to a single braid diagram v.

A crossing in v will be called *bad* if the strand that makes the overcrossing is the one whose lower endpoint is farther to the right. If v has no bad crossings, stop here.

Suppose v has at least one bad crossing. Let the *worst* crossing be a bad crossing whose undercrossing strand has lower endpoint farthest to the left. If there is more than one such bad crossing, let the worst be the one that is closest to the bottom of the diagram.

Use the skein relation to rewrite v as a linear combination of v' and v_0, where v' is the result of changing the sign of the worst crossing and v_0 is the result of removing it. Now recursively apply this procedure to v' and v_0.

Note that any bad crossings in v' and v_0 are "better" than the worst crossing of v in the sense that either the lower endpoint of their undercrossing strand is farther to the right or they have the same undercrossing strand and are closer to the top of the diagram. Thus the above algorithm must eventually terminate with a linear combination of diagrams that have no bad crossings. Any such diagram must equal T_w for some $w \in \mathfrak{S}_n$.

This algorithm shows that the T_w span $\mathcal{H}_n$. It remains to show that they are linearly independent. Note that if the algorithm is given as input a linear combination of diagrams of the form T_w, then its output will be the same linear combination. Thus it suffices to show that the output of the algorithm does not depend on the initial choice of linear combination of braid diagrams to represent a given element of $\mathcal{H}_n$. We prove this in three claims, which show that the output of the algorithm is invariant under the skein relation and Reidemeister moves of types two and three.

CLAIM 4.2. *Suppose v_+, v_- and v_0 are three braid diagrams that are identical except in a small disk where v_+ has a positive crossing, v_- has a negative crossing, and v_0 has no crossing. Then the algorithm gives the same output for both sides of the skein relation $v_+ + q_1 q_2 v_- = (q_1 + q_2) v_0$.*

PROOF. For exactly one of v_+ and v_-, the crossing inside the small disk is a bad crossing. For convenience assume it is v_+, since it makes no difference to the argument.

Suppose the worst crossing for v_+ is the crossing in the small disk. Applying the next step of the algorithm to v_+ results in a linear combination of v_- and v_0 which, by design, will exactly cancel the other two terms in the skein relation.

Now suppose the worst crossing for v_+ is not inside the small disk. Then it must be the same as the worst crossing for v_- and for v_0. Thus the next step of the algorithm has the same effect on v_+, v_- and v_0. The claim now follows by induction. $\square$

CLAIM 4.3. *If u and v are diagrams that differ by a Reidemeister move of type two then the algorithm gives the same output for u as for v.*

PROOF. As in the proof of the previous claim, we can reduce to the case where a worst crossing lies inside the small disk affected by the Reidemeister move. The claim now follows by computing the result of applying the algorithm inside the small disk. Alternatively, observe that this computation amounts to checking the case $n = 2$ of Theorem 4.1, which follows easily from the presentation of $\mathcal{H}_2$. $\square$

CLAIM 4.4. *Suppose u and v are diagrams that differ by a Reidemeister move of type three. Then the algorithm gives the same output for u as for v.*

PROOF. Once again, one solution involves a brute force computation of the algorithm. Here we describe a somewhat more comprehensible approach.

Label the three strands in the small disk in each of u and v the *front*, *back*, and *middle* strands, where the front strand makes two overcrossings, the back strand makes two undercrossings, and the middle strand makes one overcrossing and one undercrossing.

Let u' and v' be the result of changing the sign of the crossings between the front and middle strands of u and v respectively. Note that eliminating these crossings results in identical braid diagrams. Thus by Claim 4.2, the algorithm gives the same output for u as for v if and only if it gives the same output for u' as for v'.

Relabel the three strands in u' and v' so that once again the front strand makes two overcrossings, the back strand makes two undercrossings, and the middle strand makes one overcrossing and one undercrossing. Now let u'' and v'' be the output of changing the sign of the crossing between the middle and back strands of u' and v' respectively. Note that eliminating these crossings results in braid diagrams that differ by Reidemeister moves of type two. Thus by Claims 4.2 and 4.3, the algorithm gives the same output for u' as for v' if and only if it gives the same output for u'' as for v''.

We can continue in this way, alternately changing crossings between front and middle, and middle and back strands. We obtain six different versions of the Reidemeister move of type three. Each is obtained from the original by some crossing changes, and corresponds to one of the six permutations of the roles of front, middle, and back strands.

The algorithm gives the same output for u as for v if and only if it gives the same output when the relevant disks in u and v are changed to represent any one of the six versions of the Reidemeister move of type three. Thus we can choose a version to suit our convenience. In particular we can always choose the front and back strands to be the ones with lower endpoints farthest to the left and right respectively. That way there will be no bad crossings inside the small disk, and the algorithm will proceed identically for the diagrams on either side of the move.

This completes the proof of the claim, and hence of the theorem.   □

5. Trace functions

A trace function on $\mathcal{H}_n$ is a linear function $\mathrm{tr}\colon \mathcal{H}_n \to R$ such that $\mathrm{tr}(ab) = \mathrm{tr}(ba)$ for all $a, b \in \mathcal{H}_n$. Let V be the quotient of $\mathcal{H}_n$ by the vector subspace spanned by elements of the form $ab - ba$ for $a, b \in \mathcal{H}_n$. Then the trace functions of $\mathcal{H}_n$ correspond to the linear maps from V to R.

The aim of this section is to find a basis for V, and hence classify all trace functions of $\mathcal{H}_n$. This has been done by Turaev [Tur88] and independently by Hoste and Kidwell [HK90]. They actually consider a larger algebra in which the strands can have arbitrary orientations, but the result is very similar.

We define a *closed n-braid* to be a disjoint union of circles in $D \times S^1$ that intersects each disk cross-section at a total of n points. We say two closed braids are equivalent if one can be deformed to the other through a continuous family of closed braids. The *closure* of a geometric braid in $D \times I$ is the result of identifying $D \times \{0\}$ to $D \times \{1\}$. It is not difficult to show that two braids have equivalent closures if and only if they are conjugate in B_n.

Define a *diagram* of a closed braid to be a projection onto the annulus $I \times S^1$ in general position, together with crossing information, similar to the diagram of a braid. Then V is the vector space of formal linear combination of closed n-braid diagrams modulo the skein relation and Reidemeister moves of types two and three. This is an example of a *skein algebra* of the annulus.

A *partition of n* is a sequence $\lambda = (\lambda_1, \ldots, \lambda_k)$ of integers such that $\lambda_1 \geq \cdots \geq \lambda_k > 0$ and $\lambda_1 + \cdots + \lambda_k = n$. For any $m > 0$, let $b_{(m)}$ be the braid $\sigma_{m-1} \cdots \sigma_2 \sigma_1$.
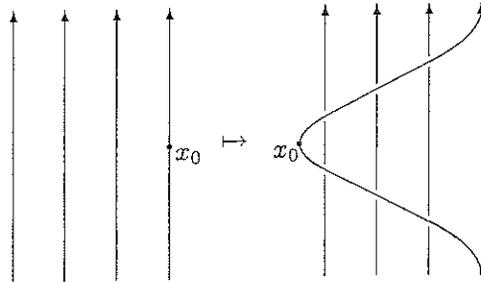
FIGURE 4. Pulling the basepoint toward $\{0\} \times S^1$

If $\lambda = (\lambda_1, \ldots, \lambda_k)$ is a partition of n, let b_λ be the braid with diagram consisting of a disjoint union of diagrams of the braids $b_{(\lambda_i)}$, in order from left to right. Let v_λ be the closure of b_λ.

THEOREM 5.1. *The set of v_λ for partitions λ of n forms a basis for V.*

We start by defining an algorithm similar to that of Theorem 4.1. There are some added complications because a strand can circle around and cross itself, and there is no "bottom" of the closed braid to use as a starting point. Therefore the first step is to choose a basepoint x_0 on the diagram that is not a crossing point. Pull x_0 in front of the other strands so that x_0 becomes the closest point to the boundary component $\{0\} \times S^1$, as suggested by Figure 4.

Consider the oriented edge that begins at x_0 and proceeds in the positive direction around the annulus. Call a crossing *bad* if this edge makes an undercrossing on the first (or only) time it passes through that crossing. If there is a bad crossing, use the skein relation to eliminate the first bad crossing the edge encounters. This process will eventually terminate with a linear combination of diagrams that have no bad crossings.

If there are no bad crossings then the loop through x_0 passes in front of every other loop in the closed braid. Furthermore, we can assume that its distance toward the front of the diagram steadily decreases as it progresses in the positive direction from x_0 until just before it closes up again at x_0. Since x_0 is the closest point to $\{0\} \times S^1$, this implies that the loop through x_0 is isotopic to $v_{(m)}$ for some positive integer m.

Isotope this loop, keeping it in front of all other strands, toward the boundary component $\{1\} \times S^1$, until its projection is disjoint from that of all other loops. Now ignore this loop and repeat the above procedure to the remainder of the closed braid diagram. This process must eventually terminate with a linear combination of closed braids of the form v_λ.

To show that the v_λ are linearly independent, it suffices to show that output is invariant under the skein relation, Reidemeister moves of types two and three, and the choices of basepoint. By induction on n we can assume that the output of the algorithm does not depend on choices of basepoint made after the first loop has been made disjoint from the other loops. Thus the algorithm produces a unique output given a diagram of a closed n-braid together with a single choice of initial basepoint x_0.

CLAIM 5.2. *The output of the algorithm is invariant under the skein relation, and under any Reidemeister move for which the basepoint does not lie in the disk affected by the move.*

PROOF. The proofs of Claims 4.2, 4.3 and 4.4 go through unchanged.   □

It remains only to prove the following.

CLAIM 5.3. *For a given diagram of a closed n-braid, the output of the algorithm does not depend on the choice of basepoint.*

PROOF. First we show that the output of the algorithm is not affected by moving the basepoint over an overcrossing. Recall that the first step of the algorithm is to pull the basepoint in front of the other strands as in Figure 4. If the basepoint is moved over an overcrossing, the output of this first step will be altered by a Reidemeister move of type two. Furthermore, the basepoint does not lie in the disk affected by this Reidemeister move. Thus the output of the algorithm is unchanged.

Now fix a diagram v of a closed n-braid. By induction, assume that the claim is true for any diagram with fewer crossings than v. By the skein relation, if v' is the result of changing the sign of one of the crossings of v then the claim is true for v if and only if it is true for v'. Thus we are free to change the signs of the crossings in v to suit our convenience.

By changing crossings and moving the basepoint past overcrossings we can move the basepoint to any other point on the same loop in v. Now suppose basepoints x_0 and y_0 lie on two distinct loops in v. We can choose each to be the closest point on its loop to the boundary component $\{0\} \times S^1$. Assume, without loss of generality, that x_0 is at least as close as to $\{0\} \times S^1$ as y_0 is. By changing the signs of crossings, we can assume that the loop through y_0 has no bad crossings. Applying the algorithm using y_0 as the basepoint then has the affect of isotoping the corresponding loop toward $\{1\} \times S^1$, keeping it in front of all other strands. This can be achieved by a sequence of Reidemeister moves of types two and three. No point in this loop is closer to $\{0\} \times S^1$ than x_0, so x_0 does not lie in the disk affected by any of these Reidemeister moves. Thus the algorithm will give the same output using the basepoint x_0 as it does using basepoint y_0.

This completes the proof of the claim, and hence of the theorem.   □

6. The Jones polynomial

The Jones polynomial is an invariant of knots and links, first defined by Jones [**Jon85**]. Jones arrived at his definition as an outgrowth of his work on operator algebras, as opposed to knot theory. To this day the topological meaning of his polynomial seems somewhat mysterious, and it has a very different flavor to classical knot invariants such as the Alexander polynomial.

After the discovery of the Jones polynomial, several people independently realized that it could be generalized to a two-variable polynomial now called the HOMFLY or HOMFLYPT polynomial. The names are acronyms of the authors of [**FYH+85**], where the polynomial was defined, and of [**PT88**], where related results were discovered independently.

The aim of this section is to show how to use the Iwahori-Hecke algebra to define a polynomial invariant of knots and links called the HOMFLY or HOMFLYPT polynomial.

Given a geometric braid b, we can obtain a closed braid in the solid torus by identifying the top and the bottom of b. Now embed this solid torus into S^3 in a standard unknotted fashion. The resulting knot or link in S^3 is called the *closure* of b. It is a classical theorem of Alexander that any knot or link in S^3 can be obtained in this way.

Let B_∞ be the disjoint union of the braid groups B_n for $n \geq 1$. For every $n \geq 1$, let

$$\iota \colon B_n \to B_{n+1}$$

be the inclusion map that adds a single straight strand to the right of any n-braid. The *Markov moves* are as follows.

- $ab \leftrightarrow ba$,
- $b \leftrightarrow \sigma_n \iota(b)$,
- $b \leftrightarrow \sigma_n^{-1} \iota(b)$,

for any $a, b \in B_n$.

THEOREM 6.1 (Markov's theorem). *Two braids have the same closure if and only if they are connected by a sequence of Markov moves.*

An R-valued link invariant is thus equivalent to a function from B_∞ to R that is invariant under the Markov moves. We now look for such a function that factors through the maps $B_n \to \mathcal{H}_n$.

Let $\iota \colon \mathcal{H}_n \to \mathcal{H}_{n+1}$ be the inclusion map $T_i \mapsto T_i$. A family of linear maps $\operatorname{tr} \colon \mathcal{H}_n \to R$ defines a link invariant if and only if it satisfies the following.

- $\operatorname{tr}(ab) = \operatorname{tr}(ba)$,
- $\operatorname{tr}(b) = \operatorname{tr}(T_n \iota(b))$,
- $\operatorname{tr}(b) = \operatorname{tr}(T_n^{-1} \iota(b))$,

for every $n \geq 1$ and $a, b \in \mathcal{H}_n$. We will call such a family of maps a *normalized Markov trace*. The usual definition of Markov trace is slightly different, but can easily be rescaled to satisfy the above conditions.

By the skein relation, the third condition on a normalized Markov trace is equivalent to

$$(1) \qquad\qquad (1 + q_1 q_2) \operatorname{tr}(b) = (q_1 + q_2) \operatorname{tr}(\iota(b)),$$

for every $n \geq 1$ and $b \in \mathcal{H}_n$. To obtain an interesting invariant, we assume from now on that $q_1 + q_2$ is a unit of R.

Let λ be a partition of n and let b_λ be the corresponding braid diagram as defined in the previous section. Let k be the number of components of b_λ, that is, the number of nonzero entries in λ. Using the Markov moves and Equation (1), it is easy to show that any normalized Markov trace must satisfy

$$(2) \qquad\qquad \operatorname{tr}(b_\lambda) = \left(\frac{1 + q_1 q_2}{q_1 + q_2} \right)^{k-1} \operatorname{tr}(\operatorname{id}_1),$$

where id_1 is the identity element of $\mathcal{H}_1$.

We now show that this equation defines a normalized Markov trace. Let b be a braid diagram. Let v be the closed braid diagram obtained by identifying the top and the bottom of b. Apply the algorithm from the previous section to write v as a linear combination of v_λ. Let $\operatorname{tr}(b)$ be as given by Equation (2) with $\operatorname{tr}(\operatorname{id}_1) = 1$.

Let v_0 be the closed braid diagram obtained by identifying the top and the bottom of $\iota(b)$. This is obtained from v by adding a disjoint loop. Now apply

the algorithm to write v_0 as a linear combination of basis elements v_λ. The added loop in v_0 remains unchanged throughout the algorithm. Thus it has the effect of adding an extra component to each term of the resulting linear combination of basis elements. This shows that tr satisfies Equation (1).

Now let v_+ be the closed braid diagram obtained by identifying the top and the bottom of $\sigma_n \iota(b)$. then v_+ is obtained from v by adding an extra "kink" in one of the strands. Now apply the algorithm to write v_+ as a linear combination of basis elements v_λ. We can assume that we never choose a basepoint that lies on the added kink. Then the added kink remains unchanged throughout the algorithm. It has no effect on the number of components in each term of the resulting linear combination of basis elements. This shows that tr is invariant under the second Markov move.

This completes the proof that tr is a normalized Markov trace. Any other normalized Markov trace must be a scalar multiple of tr. The HOMFLYPT polynomial $P_L(q_1, q_2)$ of a link L is defined to be $\text{tr}(b)$ for any braid b whose closure is L. There are many different definitions of P_L in the literature, each of which can be obtained from any other by a change of variables. They are usually specified by giving the coefficients of the three terms of the skein relation in Figure 3. As far as I know, mine is yet another addition to the collection of possible choices that appear in the literature.

The Jones polynomial V_L is given by

$$V_L(t) = P_L(-t^{\frac{1}{2}}, t^{\frac{3}{2}}).$$

If L is a knot, V_L turns out to involve only integer powers of t. This polynomial was originally defined as a trace function of the Temperley-Lieb algebra, which is a certain quotient of the Iwahori-Hecke algebra.

A somewhat tangential question is worth mentioning here. In its most open-ended form, it is as follows.

QUESTION 2. *What are the equivalence classes of braids modulo the moves*

- $ab \leftrightarrow ba$, *and*
- $b \leftrightarrow \sigma_n \iota(b)$?

In other words, what happens if the Markov move

$$b \leftrightarrow \sigma_n^{-1} \iota(b)$$

is omitted? This question was shown in [**OS03**] to be equivalent to the important problem in contact geometry of classifying transversal links up to transversal isotopy.

The *Bennequin number* of a braid $b \in B_n$ is $e - n$, where e is the sum of the exponents in a word in the generators σ_i representing b. The Bennequin number is invariant under the moves in Question 2. Thus it can be used to show that there are braids that are related by Markov moves, but not by the moves in Question 2.

Birman and Menasco [**BM**] and Etnyre and Honda [**EH**] have independently found pairs of braids that are related by Markov moves and have the same Bennequin invariant, but are not related by the moves in Question 2. Their proofs are quite complicated, and it would be nice to have a new invariant that could distinguish their pairs of braids.

7. Representations of $\mathcal{H}_n$

A *representation* of $\mathcal{H}_n$ is simply a $\mathcal{H}_n$-module. If R is a field then an *irreducible* representation of $\mathcal{H}_n$ is a nonzero $\mathcal{H}_n$-module with no nonzero proper submodules. In [**DJ86**], Dipper and James gave a complete list of the irreducible representations of $\mathcal{H}_n$. The aim of this section is to summarize their results. Our approach comes from the theory of *cellular algebras*, as defined in [**GL96**]. For convenience we will take

$$\mathcal{H}_n = \mathcal{H}_n(-1, q)$$

from now on.

Let λ be a partition of n. The *Young subgroup* $\mathfrak{S}_\lambda$ of $\mathfrak{S}_n$ is the image of the obvious embedding

$$\mathfrak{S}_{\lambda_1} \times \cdots \times \mathfrak{S}_{\lambda_k} \to \mathfrak{S}_n.$$

More precisely, it is the set of permutations of $\{1, \ldots, n\}$ that fix setwise each set of the form $\{k_i + 1, \ldots, k_i + \lambda_i\}$ where $k_i = \lambda_1 + \cdots + \lambda_{i-1}$. Let

$$m_\lambda = \sum_{w \in \mathfrak{S}_\lambda} T_w.$$

Let λ be a partition of n. Let M^λ be the left-ideal $\mathcal{H}_n m_\lambda$. We say a partition μ of n *dominates* λ if $\sum_{i=1}^j \mu_i \leq \sum_{i=1}^j \lambda_i$ for all $j \geq 1$. Let I^λ be the two-sided ideal of $\mathcal{H}_n$ generated by m_μ for all partitions μ of n that dominate λ. The *Specht module* is the quotient

$$S^\lambda = M^\lambda / (M^\lambda \cap I^\lambda).$$

Let $\operatorname{rad} S^\lambda$ be the set of $v \in S^\lambda$ such that $m_\lambda h v = 0$ for all $h \in \mathcal{H}_n$. Let

$$D^\lambda = S^\lambda / \operatorname{rad} S^\lambda.$$

THEOREM 7.1. *Suppose R is a field. Then every irreducible representation of $\mathcal{H}_n$ is of the form D^λ for some partition λ of n. If λ and μ are distinct partitions of n then D^λ and D^μ are either distinct or both zero.*

The definition of D^λ can be better motivated by defining a bilinear form on S^λ. Let $\star \colon \mathcal{H}_n \to \mathcal{H}_n$ be the antiautomorphism given by $T_w^* = T_{w^{-1}}$ for all $w \in \mathfrak{S}_n$. Note that $m_\lambda^* = m_\lambda$. The following lemma is due to Murphy [**Mur92**].

LEMMA 7.2. *If $h \in \mathcal{H}_n$ then $m_\lambda h m_\lambda = r m_\lambda$ modulo I^λ, for some $r \in R$.*

Thus we can define a bilinear form

$$\langle \cdot, \cdot \rangle \colon S^\lambda \times S^\lambda \to R$$

by

$$(h_1 m_\lambda)^*(h_2 m_\lambda) = \langle h_1 m_\lambda, h_2 m_\lambda \rangle m_\lambda.$$

Then $\operatorname{rad} S^\lambda$ is the set of $y \in S^\lambda$ such that $\langle x, y \rangle = 0$ for all $x \in \mathcal{H}_n$.

Dipper and James also determined which values of λ give a nonzero D^λ. Let e be the smallest positive integer such that $1 + q + \cdots + q^{e-1} = 0$, or infinity if there is no such integer.

THEOREM 7.3. $D^\lambda \neq 0$ *if and only if $\lambda_i - \lambda_{i+1} < e$ for all $i \geq 1$.*

Thus the work of Dipper and James completely characterizes the irreducible representations of $\mathcal{H}_n$. However, understanding these irreducible representations remains an active area of research to this day. An example of a major open-ended question in the area is the following.

QUESTION 3. *What can be said about the dimensions of D^λ?*

8. The future

My hope for the future is that the definition of B_n as a mapping class group will provide solutions to problems related to the representation theory of the Iwahori-Hecke algebra. One reason for optimism is the mechanism described in [**Big04**] to obtain representations of the $\mathcal{H}_n$ from the induced action of B_n on homology modules of configuration spaces in D_n. There I conjectured that all irreducible representations D^λ can be obtained in this way. The inner product we defined on D^λ would presumably correspond to the intersection form on homology.

Another direction for future research is to look at other quotient algebras of RB_n. After the Iwahori-Hecke algebra, the next obvious candidate is the Birman-Wenzl-Murakami algebra. A somewhat non-standard presentation of this algebra is as follows.

Let X be the following element of the braid group algebra RB_n.

$$X = q\bar{\sigma}_1 + 1 - q - \sigma_1.$$

The Birman-Wenzl-Murakami algebra is the quotient of RB_n by the following relations.

- $(q^2\sigma_1^{-1}\sigma_2^{-1} - \sigma_1\sigma_2)X = 0$,
- $(q\sigma_2^{-1} + 1 - q - \sigma_2)X = (q\sigma_1^{-1}\sigma_2^{-1} - \sigma_1\sigma_2)X$,
- $\sigma_1 X = tX$.

This algebra has an interesting history. After the discovery of the Jones polynomial, Kauffman [**Kau90**] discovered a new knot invariant which he defined directly using the knot or link diagram and a skein relation. The Birman-Wenzl-Murakami algebra was then constructed in [**BW89**], and independently in [**Mur87**], so as to give the Kauffman polynomial via a trace function. Thus the history of the Birman-Wenzl-Murakami algebra traces the history of the Jones polynomial in reverse.

QUESTION 4. *How much of this paper can be generalized to the Birman-Wenzl-Murakami algebra?*

In this direction, John Enyang [**Eny04**] has shown that the Birman-Murakami-Wenzl algebra is a cellular algebra, and used this to give a definition of its irreducible representations similar to the approach in Section 7.

Next we would like to generalize [**Big04**] to the Birman-Wenzl-Murakami algebra.

QUESTION 5. *Is there a homological definition of representations of the Birman-Wenzl-Murakami algebra?*

I believe the answer to this is yes. Furthermore, the homological construction suggests a new algebra Z_n, which would further generalize the Iwahori-Hecke and Birman-Wenzl-Murakami algebras. I will conclude this paper with a definition of Z_n and some related open questions. I hope these might be amenable to some combinatorial computations, even without the homological motivation, which is currently unclear and unpublished.

We use the notation $\sigma_{i_1\ldots i_k}$ as shorthand for $\sigma_{i_1}\ldots\sigma_{i_k}$, and $\bar{\sigma}_{i_1\ldots i_k}$ for $\sigma^{-1}_{i_1\ldots i_k}$. Define the following elements of RB_n:

$$\begin{aligned}
X_2 &= q\bar{\sigma}_1 + 1 - q - \sigma_1 \\
X_3 &= (q^2\bar{\sigma}_{21} - \sigma_{12})X_2 \\
X_4 &= (q^3\bar{\sigma}_{321} - \sigma_{123})X_3 \\
&\vdots \\
X_n &= (q^{n-1}\bar{\sigma}_{(n-1)\ldots 1} - \sigma_{1\ldots(n-1)})X_{n-1}.
\end{aligned}$$

Then Z_n is the algebra RB_n modulo the following relations:

$$\begin{aligned}
(q\bar{\sigma}_2 + 1 - q - \sigma_2)X_2 &= (q\bar{\sigma}_{21} - \sigma_{12})X_2, \\
(q^2\bar{\sigma}_{32} - \sigma_{23})X_3 &= (q^2\bar{\sigma}_{321} - \sigma_{123})X_3, \\
(q^3\bar{\sigma}_{432} - \sigma_{234})X_4 &= (q^3\bar{\sigma}_{4321} - \sigma_{1234})X_4, \\
&\vdots \\
(q^{n-1}\bar{\sigma}_{n\ldots 2} - \sigma_{2\ldots n})X_n &= (q^{n-1}\bar{\sigma}_{n\ldots 1} - \sigma_{1\ldots n})X_n.
\end{aligned}$$

Note that $\mathcal{H}_n(1, -q)$ is the quotient of Z_n by the relation $X_2 = 0$. Also the Birman-Wenzl-Murakami algebra is the quotient of Z_n by the relations $X_3 = 0$ and $\sigma_1 X_2 = tX_2$. The following basically asks if Z_n is bigger than the Birman-Wenzl-Murakami algebra.

QUESTION 6. *Does X_3 equal 0 in Z_n?*

Presumably some extra relations should be added to Z_n, such as $\sigma_1 X_2 = tX_2$, or something more general.

QUESTION 7. *What extra relations should be added to Z_n to make it finite-dimensional?*

QUESTION 8. *How much of this paper can be generalized to Z_n?*

It might be easier to first study these questions for the quotient of Z_n by the relation $X_4 = 0$.

References

[Bar03] John W. Barrett, *Geometrical measurements in three-dimensional quantum gravity*, Proceedings of the Tenth Oporto Meeting on Geometry, Topology and Physics (2001), vol. 18, 2003, pp. 97–113. MR2029691 (2005b:83034)

[BB] Joan S. Birman and Tara E. Brendle, *Braids: A survey*, Handbook of knot theory, 19–103, Elsevier B. V., Amsterdam, 2005. MR2179260

[Big99] Stephen Bigelow, *The Burau representation is not faithful for $n = 5$*, Geom. Topol. **3** (1999), 397–404 (electronic). MR1725480 (2001j:20055)

[Big02] ———, *Does the Jones polynomial detect the unknot?*, J. Knot Theory Ramifications **11** (2002), no. 4, 493–505, Knots 2000 Korea, Vol. 2 (Yongpyong). MR1915491 (2003c:57010)

[Big04] ———, *Homological representations of the Iwahori-Hecke algebra*, Geometry and Topology Monographs **7** (2004), 493–507.

[Bir74] Joan S. Birman, *Braids, links, and mapping class groups*, Princeton University Press, Princeton, N.J., 1974, Annals of Mathematics Studies, No. 82.

[BM] Joan S. Birman and William W. Menasco, *Stabilization in the braid groups-II: Transversal simplicity of knots.*

[BW89]   Joan S. Birman and Hans Wenzl, *Braids, link polynomials and a new algebra*, Trans. Amer. Math. Soc. **313** (1989), no. 1, 249–273.

[CL97]   D. Cooper and D. D. Long, *A presentation for the image of* Burau(4) $\otimes Z_2$, Invent. Math. **127** (1997), no. 3, 535–570. MR1431138 (97m:20050)

[CL98]   _____, *On the Burau representation modulo a small prime*, The Epstein birthday schrift, Geom. Topol. Monogr., vol. 1, Geom. Topol. Publ., Coventry, 1998, pp. 127–138 (electronic). MR1668343 (99k:20077)

[Dip85] Richard Dipper, *On the decomposition numbers of the finite general linear groups. II*, Trans. Amer. Math. Soc. **292** (1985), no. 1, 123–133. MR805956 (87c:20028)

[DJ86]   Richard Dipper and Gordon James, *Representations of Hecke algebras of general linear groups*, Proc. London Math. Soc. (3) **52** (1986), no. 1, 20–52. MR88b:20065

[EH]   John B. Etnyre and Ko Honda, *Cabling and transverse simplicity*, Ann. of Math. (2) **162** (2005), no. 3, 1305–1333. MR2179731

[Eny04] John Enyang, *Cellular bases for the Brauer and Birman-Murakami-Wenzl algebras*, J. Algebra **281** (2004), no. 2, 413–449. MR2098377

[FYH+85] P. Freyd, D. Yetter, J. Hoste, W. B. R. Lickorish, K. Millett, and A. Ocneanu, *A new polynomial invariant of knots and links*, Bull. Amer. Math. Soc. (N.S.) **12** (1985), no. 2, 239–246.

[GL96]   J. J. Graham and G. I. Lehrer, *Cellular algebras*, Invent. Math. **123** (1996), no. 1, 1–34. MR1376244 (97h:20016)

[HK90]   Jim Hoste and Mark E. Kidwell, *Dichromatic link invariants*, Trans. Amer. Math. Soc. **321** (1990), no. 1, 197–229. MR961623 (90m:57007)

[Jon85] Vaughan F. R. Jones, *A polynomial invariant for knots via von Neumann algebras*, Bull. Amer. Math. Soc. (N.S.) **12** (1985), no. 1, 103–111. MR86e:57006

[Kau90] Louis H. Kauffman, *An invariant of regular isotopy*, Trans. Amer. Math. Soc. **318** (1990), no. 2, 417–471. MR958895 (90g:57007)

[Lon86] D. D. Long, *A note on the normal subgroups of mapping class groups*, Math. Proc. Cambridge Philos. Soc. **99** (1986), no. 1, 79–87. MR809501 (87c:57009)

[Mur87] Jun Murakami, *The Kauffman polynomial of links and representation theory*, Osaka J. Math. **24** (1987), no. 4, 745–758.

[Mur92] G. E. Murphy, *On the representation theory of the symmetric groups and associated Hecke algebras*, J. Algebra **152** (1992), no. 2, 492–513. MR1194316 (94c:17031)

[OS03]   S. Yu. Orevkov and V. V. Shevchishin, *Markov theorem for transversal links*, J. Knot Theory Ramifications **12** (2003), no. 7, 905–913. MR2017961 (2004j:57011)

[PT88]   Józef H. Przytycki and Paweł Traczyk, *Invariants of links of Conway type*, Kobe J. Math. 4 (1988), no. 2, 115–139. MR945888 (89h:57006)

[Tur88] V. G. Turaev, *The Conway and Kauffman modules of a solid torus*, Zap. Nauchn. Sem. Leningrad. Otdel. Mat. Inst. Steklov. (LOMI) **167** (1988), no. Issled. Topol. 6, 79–89, 190. MR964255 (90f:57012)

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA AT SANTA BARBARA, CALIFORNIA 93106, USA

E-mail address: bigelow@math.ucsb.edu