

The Burau representation is faithful for $n = 4$

Vasudha Bharathram Joan S. Birman Tara E. Brendle

September 15, 2026

Abstract

In this paper we use ideas introduced earlier by Moody [20], Long [16], Long–Paton [17] and Bigelow [3] to prove the theorem of the title, that the Burau representation of the classical braid group B_4 is faithful. An immediate corollary is that the Jones representation of B_4 is also faithful.

1 Introduction

Let D_n be a disk with n marked points $p_1, p_2, \dots, p_n$ in its interior. The **braid group** B_n will be interpreted in this paper as the mapping class group $\text{Mod}(D_n)$. In 1935 Werner Burau introduced a representation of B_n into $\text{GL}_{n-1}(\mathbb{Z}[t, t^{-1}])$, now known as the reduced Burau representation [7]. We will work mainly with the unreduced Burau representation $\rho_n : B_n \rightarrow \text{GL}_n(\mathbb{Z}[t, t^{-1}])$, and henceforth we refer to ρ_n simply as the Burau representation.

Magnus–Peluso first established the faithfulness of the Burau representation for $n = 3$ by a direct algebraic calculation [22]. Moody later proved that ρ_n is not faithful for $n \geq 9$ [20]. Long–Paton built on Moody’s ideas to improve this to $n \geq 6$ [17], and Bigelow later added the case $n = 5$ [3].

In the years following Bigelow’s paper there have been many attempts to detect elements in the kernel of the remaining case ρ_4 , or to narrow the search for such, each of which contributed to our understanding of this question as viewed through the lens of various areas of mathematics; see for example Alperin–Farb–Noskov [1], Beridze–Traczyk [5, 6], Calvez–Ito [8], Datta [9], Dlugie [10], Fullarton–Shadrach [12], Gibson–Williamson–Yacobi [13], and Witzel–Zaremsky [24]. The question of faithfulness for ρ_4 also appears as Question 3.1 in Margalit’s problem list for mapping class groups [19].

Theorem 1.1 Main Theorem. *The Burau representation ρ_4 of B_4 is faithful.*

As a first step to proving this result, we will give a new and simple topological proof of the faithfulness of the Burau representation in the case $n = 3$. We will then adapt these ideas to the case of $n = 4$ by exploiting the structure of point-pushing maps in B_4 .

This approach is in contrast to earlier work on B_n for $n \geq 5$ and attempts to prove that the Burau representation ρ_4 is not faithful. In particular, we do not consider elements in the image of ρ_4 that potentially generate a free group. We refer the reader to the second author’s book [4, Theorem 3.19], where this approach was first introduced. In addition to the papers mentioned above, see also Bigelow [3, Section 3] and Moran [21] for details of other approaches.

Application to the Jones representation. In his seminal work, Jones introduced a representation of the braid group that contains the (reduced)

Burau representation as a summand; for a full definition we refer the reader to his paper [14]. As Jones points out, the faithfulness of Burau representation ρ_n for any n implies the faithfulness of the Jones representation of the braid group B_n . Therefore we immediately obtain the following corollary of our [Main Theorem](#).

Corollary 1.2 *The Jones representation of B_n is faithful for $n = 4$.*

We also refer the reader to Kasahara’s subsequent work on the Jones representation, in particular to his explanation of the equivalence between the faithfulness of the Burau representation and the Jones representation in the case $n = 4$ [15, Remark 5.6].

Strategy of proof. Let K_i denote the point-pushing subgroup of B_n obtained as the kernel of the Birman exact sequence for the disk D_n by “forgetting” the i^{th} marked point in the disk D_n ; see [Section 5](#) for more details. The intersection $\cap_{i=1}^n K_i$ of all point-pushing subgroups in B_n is known as the **Brunnian group** $Brun_n$, and it is a normal subgroup of B_n . A theorem of Long states that the Burau representation ρ_n is faithful on B_n if it is faithful on any nontrivial noncentral normal subgroup of B_n [16, Theorem 2.2]. Suppose now that $\ker(\rho_4)$ were nontrivial. By Long’s theorem, $\ker(\rho_4)$ intersects $Brun_4$ nontrivially.

Proposition 1.3 *If the Burau representation ρ_4 is faithful on its restriction to the Brunnian group $Brun_4$, then it is faithful on B_4 .*

Building on ideas of Long–Paton and Bigelow, we associate to each braid in B_n a sequence of disks in D_n that carry certain combinatorial data. We use this technology to give a new proof of Magnus–Peluso’s theorem that the Burau representation ρ_3 is faithful [22]. Our proof yields the more general result that any n -strand braid whose associated disk sequence satisfies a certain **parity condition**, defined in [Section 4](#), does not lie in the kernel of ρ_n .

Braids in B_n do not in general satisfy this parity condition when $n \geq 4$. However, we show that when $n = 4$, the parity condition “almost” holds for any point-pushing braid $\Phi \in K_4$ that admits a factorization as a **proper product** of certain push-maps. The next step is to “correct” Φ by embedding K_4 in B_5 , that is, we use Φ to construct a 5-braid satisfying the parity condition. We then apply a result of Moody to conclude that the original braid $\Phi \in K_4$ does not lie in the kernel of ρ_4 . The final step is to show that any braid in K_4 is conjugate to another braid in K_4 that can be realized as a proper product of push-maps.

Finally, we remark that, *a priori*, it is not necessarily clear that $Brun_n$ is nontrivial. The simplest example of a Brunnian braid is a three-strand braid that closes to form the Borromean rings. [Figure 1.1](#) gives an example of a nontrivial element of $Brun_4$, the case of interest to us, and a similar construction yields nontrivial examples of Brunnian braids in B_n for any n . Moreover, Whittlesey has shown that all nontrivial Brunnian braids are pseudo-Anosov [23]. We note that Whittlesey’s results are stated for the mapping class group of the $(n+1)$ -punctured sphere, but the result holds in B_n ; see Lee–Song’s discussion of this point [18, Section 1].




Figure 1.1 A Brunnian 4-braid β .

Kernel and image of Burau. There is a great deal of literature on the faithfulness of the Burau representation, and this paper settles that question in the final outstanding case. The questions of giving both a useful characterization of its kernel and of its image remain wide open in general, and the authors believe that the methods of this paper may be useful in achieving progress on these important questions.

Outline of the paper. In [Section 2](#), we introduce the Moody polynomial of a braid, an invariant that provides an obstruction for an element to lie in the kernel of the Burau representation. Building on Moody’s ideas and their development by Long–Paton, Bigelow later introduced combinatorial tools that are useful for computing the obstruction; we describe these in [Section 3](#). Using these tools, in [Section 4](#), we give our new topological proof of faithfulness of the Burau representation in the case $n = 3$. We introduce our key technical tool, **proper products** of point-pushing braids, in [Section 5](#). In [Section 6](#) we construct point-pushing maps arising as particular proper products, and use these to establish that the Burau representation ρ_4 is faithful on its restriction to the point-pushing group Brun_4 ; by [Proposition 1.3](#), this completes the proof of the [Main Theorem](#). Finally, we give an example illustrating the constructions used in our proof in [Section 7](#).

Acknowledgements. The authors heartily thank Dan Margalit for many helpful discussions and are especially grateful for his extensive and valuable feedback on initial drafts of this paper. The second author would also like to thank him for his interest in this project from its inception seven years ago, when the first two authors had begun to work on it together. The authors also thank Khalid Bou-Rabee, Michael Dougherty, and Benson Farb for additional helpful comments on earlier drafts and Ken Birman for further helpful remarks. The first author thanks David Gabai for his invaluable mentorship and constant encouragement. The third author is also grateful to the Institute for Computational and Experimental Research in Mathematics (ICERM) at Brown University, where in 2022 she was a participant in the semester-long program “Braids”, during which initial discussions that led to her joining this project took place, and to Muffy Calder for her support and encouragement throughout. AI was used to proofread a draft of the manuscript.

2 The Moody polynomial and winding numbers

While Burau originally defined his representation by giving its values on the standard Artin generators of B_n (see, for example, [\[2, Section 4.2\]](#)), it will be more useful for present purposes to define ρ_n via the action of B_n on the disk D_n , which we consider here as an n -times punctured disk. Let p_* denote a point in the boundary of the disk D_n , and let $x_1, \dots, x_n$ denote the standard free generators of $\pi_1(D_n, p_*)$. We consider the mapping $\pi_1(D_n, p_*) \rightarrow \mathbb{Z}$ that

takes a word in $x_1, \dots, x_n$ to the sum of its exponents, and let $\widetilde{D}_n$ denote the covering space associated with its kernel.

The group of covering transformations of $\widetilde{D}_n$ is isomorphic to $\mathbb{Z}$, which we denote as a multiplicative group generated by t . Following the treatment of Long–Paton, it will be convenient for us to work with the relative homology group $H_1(\widetilde{D}_n, \{\tilde{p}_*\})$, where $\{\tilde{p}_*\}$ denotes the full pre-image of the basepoint p_* in the cover $\widetilde{D}_n$. The braid group B_n acts on the $\mathbb{Z}[t, t^{-1}]$ -module $H_1(\widetilde{D}_n, \{\tilde{p}_*\})$, which is free of rank n ; this action is the **unreduced Burau representation** ρ_n , which we will refer to simply as the Burau representation.

We next record a basic fact that follows directly from the definition of ρ_n .

Observation 2.1 Let $f : B_n \rightarrow B_{n+1}$ be the inclusion map corresponding to the standard embedding of D_n into D_{n+1} . If $\Phi \in B_n$ lies in the kernel of ρ_n , then $f(\Phi)$ lies in the kernel of ρ_{n+1} .

Observation 2.1 will be crucial in our proof of the [Main Theorem](#).

Moody polynomials. In what follows, we will closely follow Bigelow’s notation and exposition [3]. We will use the term **arc** to refer to a proper embedding of an interval in the disk D_n and **subarc** to refer to its restriction to a compact connected subinterval. Our arcs will normally be oriented, and it will often be convenient to consider arcs in terms of their image in the disk D_n rather than as functions per se.

Let α and β be two oriented arcs in D_n , considered as a disk with marked points, whose endpoints lie in the set $\{p_1, \dots, p_n, p_*\}$, and choose two corresponding lifts $\tilde{\alpha}$ and $\tilde{\beta}$ in $\widetilde{D}_n$. We define the **Moody polynomial** $\mathbb{M}(\alpha, \beta) \in \mathbb{Z}[t, t^{-1}]$ **of the oriented arcs α and β** as follows:

$$\mathbb{M}(\alpha, \beta) = \sum_{\ell \in \mathbb{Z}} (t^\ell \cdot \tilde{\alpha}, \tilde{\beta}) t^\ell$$

where $(t^\ell \cdot \tilde{\alpha}, \tilde{\beta})$ denotes the algebraic intersection number of the lift $t^\ell \cdot \tilde{\alpha}$ with $\tilde{\beta}$ in the cover $\widetilde{D}_n$. We note that the algebraic intersection number of arcs does not include intersections at any shared endpoints; in our case, we will always choose our arcs so that the endpoints of α are disjoint from those of β . We also note that the Moody polynomial $\mathbb{M}(\alpha, \beta)$ is only well defined up to multiplication by a power of t , due to its dependence on our choice of the lifts $\tilde{\alpha}$ and $\tilde{\beta}$. Our convention will be to choose our lifts so that the first point at which β crosses α is assigned the monomial $\pm t^0$.




Figure 2.1 The zeroth, first, and second “decks” in the universal cyclic cover $\widetilde{D}_4$ of the disk D_4 .

Following Bigelow [3] and Long–Paton [17], we now describe a planar interpretation of the Moody polynomial that enables us to work entirely in the base space D_n rather than in the cover. Without loss of generality, we can assume that the oriented arcs α and β intersect transversely in finitely many

points; we label these $q_1, \dots, q_m$ according to the order in which they appear as we traverse the arc β according to its orientation. Each intersection point q_i corresponds to a single point of intersection between two lifts $\tilde{\alpha}$ and $\tilde{\beta}$, and hence each point q_i contributes a monomial $\pm t^{k_i}$ to the overall sum in the Moody polynomial. Here the exponent k_i is such that $\tilde{\beta}$ and $t^{k_i} \cdot \tilde{\alpha}$ cross at a lift of q_i , and the sign ϵ_i of the monomial is the sign of that crossing where $\epsilon_i \in \{-1, 1\}$. Then an equivalent formulation of the Moody polynomial of the arcs α and β is as follows:

$$\mathbb{M}(\alpha, \beta) = \sum_{i=1}^m \epsilon_i t^{k_i}.$$

We emphasize that it may happen that $k_i = k_j$ for $i \neq j$. Indeed, our business in this paper will be to determine conditions under which this could happen. When we write the Moody polynomial as the sum of m terms of the form $\epsilon_i t^{k_i}$ without combining like terms, we will refer to this as the **unsimplified Moody polynomial**; when we combine all like terms we will refer to this as the **simplified Moody polynomial**. Moody showed that the polynomial $\mathbb{M}(\alpha, \beta)$ completely encodes the faithfulness of the Burau representation [20]. We will use one direction of his characterization; see [Theorem 2.3](#).

Conventions and color-coding. For the remainder of this paper, we will always choose the arc α to be the horizontal arc joining p_1 to p_2 oriented from left to right, and we will color it blue in all figures, as shown in [Figure 2.2](#). We also let β_*^3 denote the oriented arc shown in [Figure 2.2](#) from the basepoint p_* to the puncture p_3 and color it red.




Figure 2.2 The blue arc α joins p_1 to p_2 . The red arc β_*^3 joins the basepoint on ∂D to the point p_3 .

When we apply a braid $\Phi \in B_n$ to the arc β_*^3 , we will usually denote this by $\beta = (\beta_*^3)\Phi$ and also color β red; note that we write Φ on the right of the arc to which we are applying it because we will follow the tradition of writing multiplication of braids from left to right.

Furthermore, from this point on, a blue (respectively red) arc will always be a subarc of α (respectively β). Later we will add gold to our list of special colors. Our sign convention will be to say that q_i is a positive crossing, that is, $\epsilon_i = 1$, if β is directed downwards at q_i and that q_i is a negative crossing with $\epsilon_i = -1$ if β is directed upwards at q_i .

For any arcs γ and δ in the disk D_n , we let $\iota(\gamma, \delta)$ denote their geometric intersection number. Consider now our fixed choice of arc α . A remarkable result of Moody states that the non-faithfulness of the Burau representation is equivalent to the existence of an arc β joining the basepoint p_* to the marked point p_3 in D_n such that $\iota(\alpha, \beta) > 0$ and $\mathbb{M}(\alpha, \beta) = 0$ [20]. With this in mind,

we define the **Moody polynomial of the braid** $\Phi \in B_n$ as follows:

$$\mathbb{M}_\Phi = \mathbb{M}(\alpha, (\beta_*^3)\Phi).$$

We remark that by taking the boundary of a regular neighborhood of the arc $(\beta_*^3)\Phi$, we can view as a loop based at p_* traveling around p_3 , and that $\mathbb{M}_\Phi$ is well defined on the homology class of a lift of $[(\beta_*^3)\Phi]$ in $H_1(\widetilde{D}_n, \{\widetilde{p}_*\})$, which we denote by $[(\beta_*^3)\Phi]$. We also emphasize that an element $\Phi \in B_n$ uniquely determines an arc $\beta = (\beta_*^3)\Phi$, but the converse does not hold: there exist infinitely many pairs $\Phi_1, \Phi_2 \in B_n$ for which $(\beta_*^3)\Phi_1 = (\beta_*^3)\Phi_2$. We refer the reader to Long–Paton [17, Section 1] and Bigelow [3, Theorem 1.4] for further details. For our purposes, we do not need the full strength of Moody’s result; we will only require one direction of the equivalence.

Theorem 2.2 (Moody) *Let $n \geq 3$. If, given any oriented arc β from p_* to p_3 such that $\iota(\alpha, \beta) > 0$, we have that $\mathbb{M}(\alpha, \beta) \neq 0$, then the Burau representation ρ_n is faithful. Equivalently, if we have that $\mathbb{M}_\Phi \neq 0$, for any braid Φ such that $\Phi(p_3) = p_3$ and $\iota(\alpha, (\beta_*^3)\Phi) > 0$, then the Burau representation ρ_n is faithful.*

The equivalence of the two statements follows from the fact that every oriented arc β from p_* to p_3 arises as $\beta = (\beta_*^3)\Phi$ for some $\Phi \in B_n$ by the change of coordinates principle; note that such a braid Φ necessarily fixes the point p_3 . It will be useful for our purposes to adapt Moody’s theorem to obtain a specific obstruction to an individual braid lying in the kernel of ρ_n , for which we give a short proof for the sake of completeness.

Theorem 2.3 (Moody) *Let $\Phi \in B_n$ for $n \geq 3$. If $\mathbb{M}_{\Phi \cdot \Gamma} \neq \mathbb{M}_\Gamma$ for some $\Gamma \in B_n$, then Φ does not lie in the kernel of ρ_n .*

Proof. Suppose that Φ lies in the kernel of ρ_n . Then $\Gamma^{-1} \cdot \Phi \cdot \Gamma$ must also lie in the kernel. We have that $(\beta_*^3) \cdot \Gamma$ corresponds to an element $[(\beta_*^3) \cdot \Gamma]$ in the homology of the cover $\widetilde{D}_n$ relative to the pre-image of a basepoint. Since $\Gamma^{-1} \cdot \Phi \cdot \Gamma$ acts trivially on this homology group, we have that $[(\beta_*^3) \cdot (\Phi \cdot \Gamma)] = [((\beta_*^3) \cdot \Gamma) \cdot (\Gamma^{-1} \cdot \Phi \cdot \Gamma)] = [(\beta_*^3) \cdot \Gamma]$. The result now follows from our previous observation that $\mathbb{M}_\Phi$ well defined on the homology class $[(\beta_*^3)\Phi]$ in $H_1(\widetilde{D}_n, \{\widetilde{p}_*\})$. ■

In Section 6 we will identify certain products of braids in B_4 to which we will then apply Theorem 2.3 in order to establish the faithfulness of ρ_4 . Moreover, as discussed in Section 1, we ultimately reduce our search to the Brunnian subgroup Brun_n in B_n , where all nontrivial braids are pseudo-Anosov, and hence for all $\Phi \in \text{Brun}_n$ we have that $\iota(\alpha, (\beta_*^3) \cdot \Phi) > 0$. We will also use the following special case of Theorem 2.3.

Corollary 2.4 *Let $\Phi \in B_n$ for $n \geq 3$. If $\mathbb{M}_\Phi \neq 0$, then Φ does not lie in the kernel of ρ_n .*


3 Disk sequences and total winding number

Let $\beta = (\beta_*^3)\Phi$, and as above, we label the points of intersection of α with β by $q_1, \dots, q_m$ according to the order in which we encounter them as we travel along β . For $i = 1, \dots, m-1$, we let α_i and β_i denote the subarcs of α and β , respectively, that join the intersection point q_i to q_{i+1} . Note that, with our choice of notation, this means that each arc β_i contains no intersection point q_j in its interior, while an arc α_i may contain a number of other intersection points in its interior.

By construction, for each $i = 1, \dots, m-1$, we have that $\alpha_i \cup \beta_i$ bounds a k -punctured disk Δ_i . We define the **total winding number** W_i associated to the disk Δ_i to be equal to k if β_i is oriented clockwise with respect to Δ_i , or to be equal to $-k$ otherwise. We will refer to the sequence $W_1, \dots, W_{m-1}$ as the **winding number sequence** of Φ , and to the sequence $\Delta_1, \Delta_2, \dots, \Delta_{m-1}$ as the **disk sequence** of Φ ; see Figure 3.1 for examples.

We will consider two such disks Δ_i, Δ_j to be **equivalent** if there is an isotopy of (D_n, α) taking the β -component of $\partial\Delta_i$ to the β -component of $\partial\Delta_j$; in other words α is fixed setwise and the endpoints of the β -arc must be contained in α at each level of the isotopy). We note that this notion of equivalence corresponds to allowing isotopies of the arc β in the disk D_n .

We will see that the combinatorial data of the winding number sequence is sufficient for determining the faithfulness of the Burau representation in the case $n = 3$, and that the more detailed information of the disk sequence is required in the case $n = 4$. The following lemma appears as a remark in Bigelow's paper [3, Section 3]. His remark describes the relationship between winding numbers and the degree k_i of the monomial that each intersection point q_i contributes to the Moody polynomial.



(a) The arcs α and β give rise to a singleton winding number sequence given by $W_1 = 1$, with corresponding Moody polynomial $1 - t$.

(b) The winding number sequence for the arcs is given by $1, 1, -1, -2, 1$; its Moody polynomial is $2 - 2t + t^2 - t^{-1}$.

Figure 3.1

Lemma 3.1 (Bigelow, [3]) *Let q_i and q_{i+1} be two points in $\alpha \cap \beta$ that are consecutive with respect to the oriented arc β , and let k_i and k_{i+1} be the exponents of the monomials corresponding to q_i and q_{i+1} . If W_i denotes the i^{th} term in the winding number sequence for β , then we have:*

$$W_i = k_{i+1} - k_i. \quad (1)$$

We record here the following useful corollary of Lemma 3.1, also observed by Bigelow:

Corollary 3.2 *Let $W_1, \dots, W_m$ be a winding number sequence for a braid in B_n , and let k_i and k_j be the exponents of the Moody monomials corresponding to two points q_i and q_j in $\alpha \cap \beta$. Then $k_i = k_j$ if and only if $\sum_{k=i}^{j-1} W_k = 0$.*

Equipped with this useful criterion for tracking repeats of an exponent occurring in the Moody polynomial, we will proceed in the next section to consider the question of faithfulness of ρ_n when $n = 3$.

4 Faithfulness for three strands

The following theorem was first proved by Magnus and Peluso in 1969 using purely algebraic methods. We give a new proof.

Theorem 4.1 (Magnus–Peluso, [22]) *The Burau representation ρ_3 is faithful.*

Proof. Using the same notation as in the previous section, we let $\Phi \in B_3$ be an element such that the geometric intersection of α and $\beta = (\beta_*^3)\Phi$ is equal to $m > 0$. By [Theorem 2.2](#), it suffices to show that any single term of the simplified Moody polynomial $\mathbb{M}_\Phi = \sum_{i=1}^m \epsilon_i t^{k_i}$ is nonzero. If $m = 1$, then the Moody polynomial consists of a single linear term, and hence is nonzero. Suppose now that $m > 1$, and let $W_1, \dots, W_{m-1}$ denote the winding number sequence of the braid Φ . Each corresponding disk is bounded by a subarc of α and a subarc of β and contains 1, 2, or 3 marked points. The assumption that α and β are in minimal position then implies that, up to reflection about the coordinate axes and up to homeomorphism of the disk preserving each marked point p_i pointwise and the arc α setwise, each disk in the disk sequence is equivalent to one of the three types of regions illustrated in [Figure 4.1](#), possibly after applying a symmetry.




Figure 4.1 Possible disks in the disk sequence of an element of B_3 .

Notice that for each i such that $|W_i| = 2$, the sign ϵ_i of the two points of intersection of β with α must be the same, that is, $\epsilon_i = \epsilon_{i+1}$. Furthermore, if $|W_i| \in \{1, 3\}$, then we have that $\epsilon_i = -\epsilon_{i+1}$.

Suppose now that $k_i = k_j$ for two distinct terms in the unsimplified Moody polynomial. Then $\sum_{\ell=i}^{j-1} W_\ell = 0$ by [Corollary 3.2](#). Hence for $i \leq \ell \leq j - 1$ we must have an even number of terms W_ℓ with $|W_\ell| \in \{1, 3\}$. By our observations in the previous paragraph, it follows that $\epsilon_i = \epsilon_j$. In other words, no cancellation among terms of degree k_i can occur in the summation above. It follows that the Moody polynomial of Φ is not equal to zero, and therefore the Burau representation of B_3 is faithful. $\blacksquare$

The parity condition. We note that the only information that we used in the proof of [Theorem 4.1](#) was the fact that disks containing an odd number of punctures correspond to a change of sign of corresponding coefficients in the Moody polynomial, while disks containing an even number of punctures preserve the sign of the corresponding coefficients. We will record this key observation as a lemma.

To make this precise, let q_i be a point of intersection of the oriented arcs α and β . As in [Section 2](#), we let ϵ_i denote the algebraic intersection number $\hat{i}(\beta, \alpha)$ at q_i . If Δ is a disk in the sequence for Φ corresponding to the pair of points $q_i, q_{i+1} \in \alpha \cap \beta \in \partial\Delta$, then we will say that Δ is **sign-changing** if $\epsilon_i = -\epsilon_{i+1}$ and **sign-preserving** otherwise. Let $\mathcal{P}(\Delta)$ denote the set of marked points in D_n that are contained in the interior of Δ . We will say that a braid Φ satisfies the **parity condition** if a disk Δ in its disk sequence is sign-changing if and only if $|\mathcal{P}(\Delta)|$ is odd.

In other words, our proof of [Theorem 4.1](#) reduces to showing that every 3-strand braid satisfies the parity condition. Our next lemma shows that the parity condition gives a sufficient criterion for a braid $\Phi \in B_n$ to not be contained in the kernel of the Burau representation ρ_n .

Recall from [Section 2](#) that the Moody polynomial of a braid ϕ is given by

$$\mathbb{M}_\phi = \sum_{i=1}^m \epsilon_i t^{k_i}$$

where each $\epsilon_i t^{k_i}$ is the monomial associated to the i -th crossing of α and $\beta = (\beta_*^3)\phi$, and where $m = \iota(\alpha, \beta)$. We say that the Moody polynomial of ϕ **admits no cancellations** if there is no pair of indices i, j such that $k_i = k_j$ and $\epsilon_i = -\epsilon_j$ in the above expression.

Lemma 4.2 *Let $\phi \in B_n$. If ϕ satisfies the parity condition, then the Moody polynomial of ϕ admits no cancellations.*

Proof. By [Corollary 3.2](#), if $k_i = k_j$ then $|\mathcal{P}(\Delta_i)| + \dots + |\mathcal{P}(\Delta_{j-1})| = 0$. In particular this sum is even. However, by the assumption that a given disk in the disk sequence is sign-changing if and only if it contains an odd number of punctures, this means there are an even number of sign-changing disks in the subsequence $\Delta_i, \dots, \Delta_{j-1}$. In particular we have $\epsilon_i = \epsilon_j$ and so this completes the proof. $\blacksquare$

In general, n -strand braids do not satisfy the parity condition when $n \geq 4$. Our strategy for proving the [Main Theorem](#) will be to identify certain 4-braids satisfying the parity condition.

5 Push-maps and minimal position

As discussed in [Section 1](#), in order to prove our [Main Theorem](#), it suffices to establish that ρ_4 is faithful when restricted to the subgroup of B_4 consisting of point-pushing maps, also known simply as push-maps. In light of [Lemma 4.2](#), our strategy will be to identify elements of this subgroup satisfying the parity condition. In this section, as a first step in that direction, we will characterize configurations of the various arcs and curves involved that lead to the creation of bigons between α and the image of β_*^3 under the action of push-maps.

For the purpose of defining push-maps, it is convenient to use the notation $\text{Mod}(D_n)$ to denote the braid group B_n . Let $p \in D_n$ be one of the marked points in D_n , and let $\text{Mod}(D_n, p)$ denote the subgroup of $\text{Mod}(D_n)$ that fixes p . There is a forgetful map $\text{Mod}(D_n, p) \rightarrow \text{Mod}(D_{n-1})$, and the Birman exact sequence for D_n identifies the kernel of this map with $\pi_1(D_{n-1}, p)$:

$$1 \rightarrow \pi_1(D_{n-1}, p) \xrightarrow{\text{Push}} \text{Mod}(D_n, p) \xrightarrow{\text{Forget}} \text{Mod}(D_{n-1}) \rightarrow 1$$

Given a loop Γ in $\pi_1(D_{n-1}, p)$, we can consider the **push-map** obtained by pushing the point p along Γ , and the image of $\pi_1(D_{n-1}, p)$ in $\text{Mod}(D_n, p) \subset \text{Mod}(D_n)$ is an example of a **point-pushing subgroup**. We refer the reader to Farb-Margalit's book for a detailed discussion of point-pushing maps and subgroups [\[11, Section 4.2.2\]](#).

Now, let $K_i \cong \pi_1(D_{n-1}, p_i)$ denote the point-pushing subgroup of B_n corresponding to the marked point p_i in the disk D_n ; see [Figure 5.1](#) for an example of an element of the point-pushing group K_4 in B_4 and its effect on the arc β_*^3 . We remind the reader that throughout the paper, our figures show the arc α in blue and arcs of the form $\beta = \Phi(\beta_*^3)$ in red. Henceforth we will also show any loops representing push-maps in gold.




Figure 5.1 The gold loop represents an element of the subgroup K_4 . The image β of β_*^3 under the push-map around the gold loop is shown in red on the right.

In this part of the paper, we do not need to restrict to the case $n = 4$, and so we will work in the disk D_n , the braid group B_n , and its subgroup K_n for $n \geq 4$. Before we begin, we note the following.

- Without loss of generality, we can always choose a representative of a loop Γ that intersects itself minimally and transversely, and hence there are finitely many points of self-intersection.
- The results in this section hold if we replace K_n with any other point-pushing subgroup K_j where $j \geq 4$.
- In this section we do not assume that our push-maps necessarily correspond to simple loops. However, we only require simple loops in order to prove [Theorem 6.9](#) below.

Bigon-forming polygons. Let $\Gamma \in K_n$, and let β be any simple arc joining p_* to p_3 . We wish to identify local configurations of subarcs of α, β , and Γ that lead to the creation of a bigon between $(\beta)\Gamma$ and α . As a warm-up, suppose that as we travel along Γ there is a subarc of β that precedes a bigon between Γ and α . Then pushing along Γ will create a bigon (in fact, two); see [Figure 5.2\(a\)](#). Of course, in general we can avoid such a scenario simply by assuming that all curves are in pairwise minimal position. However, this points us to a more general kind of phenomenon in the case where Γ is not simple: a sequence of subarcs of Γ that, taken together, effectively form a bigon with Γ , as shown in [Figure 5.2\(b\)](#). Then, depending on the ordering of these subarcs (as determined by the orientation of Γ), pushing along Γ can similarly create a bigon.




Figure 5.2

To make this more precise, given two subarcs $\Gamma^{(1)}, \Gamma^{(2)}$ of Γ in D_n with disjoint interiors, we will say that $\Gamma^{(1)}$ **precedes** $\Gamma^{(2)}$ and write $\Gamma^{(1)} < \Gamma^{(2)}$ if the initial point of $\Gamma^{(1)}$ precedes the initial point of $\Gamma^{(2)}$ as we trace out Γ . We will further refer to a union of subarcs $\Gamma^{(1)}, \dots, \Gamma^{(\ell)}$ that are pairwise disjoint as a **piecewise Γ -arc of length ℓ** if the following conditions are satisfied:

- for each $i \in \{1, \dots, \ell - 1\}$, the final point of $\Gamma^{(i)}$ is a self-intersection point of Γ in D_n , such that the final point of $\Gamma^{(i)}$ is the initial point of $\Gamma^{(i+1)}$;
- $\Gamma^{(1)} < \Gamma^{(2)} < \dots < \Gamma^{(\ell)}$; and

- the union $\Gamma^{(1)} \cup \dots \cup \Gamma^{(\ell)}$ in D_n contains no loops.

We will now identify three types of local configurations (in addition to bigons) between α , β , and Γ that lead to the creation of bigons between α and $(\beta)\Gamma$.

1. *Generalized bigons.*

Consider now a piecewise Γ -arc γ of length ℓ where the initial point of $\Gamma^{(1)}$ and the final point of $\Gamma^{(\ell)}$ are both points of intersection of Γ with α , so that γ together with the α -subarc with these two endpoints forms a $(\ell + 1)$ -gon in D_n containing no marked points in its interior. We will refer to such an $(\ell + 1)$ -gon as a **generalized bigon**. Figure 5.2(b) shows a generalized bigon with $\ell = 2$. If the piecewise Γ -arc in a generalized bigon is preceded by a point of intersection with β , then $(\beta)\Gamma$ and α will form a bigon in the same way as if the piecewise Γ -arc were replaced by a single subarc of Γ . We emphasize that the orientation and ordering of the Γ -subarcs as we traverse Γ is crucial here, since otherwise Γ would not push the β -subarc all the way around to form a bigon between $(\beta)\Gamma$ and α . Finally, a “standard” bigon is a generalized bigon with a piecewise Γ -arc γ of length 1.

2. *Generalized trigons.*

Figure 5.3(a) shows another configuration of curves that leads to the creation of a bigon, namely a trigon in D_4 whose boundary consists of a subarc from each of the arcs α and β , and the loop Γ , where the Γ -subarc is oriented from the β -arc to the α -arc. We will refer to such a configuration of curves as a **β -to- α trigon**. The orientation of the Γ -subarc is crucial here: if its orientation were reversed, then $(\beta)\Gamma$ does not form any bigons with α in this region. As with bigons, we can replace the single subarc of Γ here with a piecewise Γ -arc of any length; we refer to any such configuration as a **generalized β -to- α trigon**. See Figure 5.3(b) for an example.





Figure 5.3

3. *Generalized rectangles.*

Another configuration of curves that may create bigons between α and $(\beta)\Gamma$ arises when two parallel subarcs of Γ traveling between β and α form a rectangle containing no marked points; we will call this a **β -to- α rectangle**. As with bigons and trigons, we can use two piecewise Γ -arcs instead of two Γ -arcs, and we refer to any such configuration as a **generalized rectangle**; see Figure 5.4(a) for an example. We emphasize that in the case of a generalized rectangle, we do not necessarily require the two piecewise Γ -arcs to be disjoint. In other words, we allow “degenerate” generalized rectangles, in which some number of the Γ -segments

involved can play a role in both piecewise Γ -arcs; we will see examples of this in the proof of the next lemma (see Figure 5.4(b)). Note also that every generalized rectangle includes two crossings of Γ and α with the same sign that are adjacent, as the crossings are ordered along α .



(a) An example of a generalized rectangle with piecewise Γ -arcs of length 1 and length 2.

(b) A degenerate generalized rectangle formed by $\Gamma^{(1)} \cup \Gamma^{(2)} \cup \Gamma^{(3)}$ and $\Gamma^{(1)} \cup \Gamma^{(2)} \cup \Gamma^{(4)}$.

Figure 5.4

4. *Generalized bigons/trigons/rectangles with basepoint.*

Let Γ_I denote the subarc of Γ from p_4 to its first point of intersection with α , and similarly let Γ_F denote the subarc of Γ from its final point of intersection with α back to p_4 . We next consider a configuration that would become a generalized polygon of any of the previous three types upon the deletion of the puncture p_4 . If such a configuration also has one additional subarc γ of Γ from a β -subarc of the generalized polygon to α , satisfying the following three properties:

- (a) γ is the union of a subarc of Γ_I together with a subarc of Γ_F (including the basepoint p_4);
- (b) γ does not cross α or β except at its endpoints;
- (c) Γ_I crosses α ; and
- (d) Γ_F crosses β .

Such a configuration also results in a bigon after applying Γ to the arc β ; see Figure 5.5 for an example.




Figure 5.5 A generalized rectangle with basepoint.

In the proof of the next lemma, we will see that generalized bigons, β -to- α trigons, and rectangles, possibly with basepoint, are the only types of configurations that give rise to bigons between $(\beta)\Gamma$ and α . Hence we will refer to these configurations collectively as **bigon-forming polygons**.

Proper products. Let $\Phi \in K_n$, and let $\beta = (\beta_*^3)\Phi$. Our next step will be to establish a set of conditions on a loop Γ based at p_n that will imply $(\beta)\Gamma$ and α are in minimal position. To that end, we let Γ_I denote the **initial component** of Γ , that is, the subarc of Γ from p_n to the first crossing of Γ with α . Similarly, we let Γ_F denote the **final component** of Γ , traveling from its final crossing with α back to p_n . We also let β' denote the β -component of the innermost disk in the disk sequence of Φ that contains only p_n if such a disk exists; otherwise we set $\beta' = \emptyset$. We say that $\Phi \cdot \Gamma$ is a **proper product** if we can choose a representative of the loop Γ that is in minimal position with respect to α and β , and satisfying the following two conditions:

1. there are no bigon-forming polygons between α , β , and Γ ; and
2. one of the following two conditions holds:
 - (a) $\Gamma_I \cap \beta' \neq \emptyset$; or
 - (b) $\Gamma_I \cap \beta' = \emptyset$ and $\Gamma_F \cap \beta' = \emptyset$.

The second condition ensures that postcomposing with Γ does not “undo” the creation (by Φ) of a disk containing only p_4 . Figure 5.6 depicts the local picture for each type of configuration: in first case (on the left in Figure 5.6), reversing the orientation of Γ has the effect of swapping the roles of Γ_I and Γ_F , resulting in a loop that does not form a proper product. In practice, we will generally be given a push-map Φ , and then seek the second factor Γ to form a proper product with Φ . We will also refer to $\Phi \cdot \Gamma_1 \cdot \Gamma_2$ as a proper product if both $\Phi \cdot \Gamma_1$ and $(\Phi \cdot \Gamma_1) \cdot \Gamma_2$ are proper products.




Figure 5.6 Two examples of the local picture for Condition 2 of a proper product $\Phi \cdot \Gamma$ indicating the initial and final components of Γ in each case; reversing the orientation of Γ in the left-hand figure gives a non-example.

The next lemma gives a straightforward criterion for ensuring that the product of two push-maps does not yield any bigons between α and the image of β_*^3 under the product.

Proposition 5.1 *Let $\Phi \in K_n$, and let $\beta = \Phi(\beta_*^3)$. If Γ is a loop based at p_n such that $\Phi \cdot \Gamma$ is a proper product, then the arcs $(\beta)\Gamma$ and α are in minimal position.*

Proof. We begin by choosing a specific representative of Γ so that α, β and Γ are pairwise in minimal position with no triple points of intersection. Up to isotopy we can assume that $(\beta)\Gamma$ and α are transverse in D_n . Consider now a point $s \in (\beta)\Gamma \cap \alpha$. If s also lies in $\beta \cap \alpha$ we will refer to this as an **original** point of intersection; otherwise we say that s is a **new** point of intersection. For the sake of simplicity, in what follows we will refer only generalized bigons,

trigons, and rectangles; the proof goes through when adding in each case with basepoint.

Suppose now that α and $(\beta)\Gamma$ form a bigon, and suppose further that the two vertices of this bigon are both new points of intersection. New points of intersection occur when a β -subarc precedes an α -subarc as we travel along Γ . Therefore our two vertices correspond to adjacent crossings of Γ and α of the same sign or of different signs, as shown in Figure 5.7. If they are the same sign, this implies the existence of a generalized rectangle between Γ , α , and β , which contradicts our assumption that $\Phi \cdot \Gamma$ is a proper product. If they are different signs, this implies the existence of a generalized bigon between Γ and α , which again contradicts our assumption that $\Phi \cdot \Gamma$ is a proper product. Therefore at least one vertex of any bigon between α and $(\beta)\Gamma$ must be an original point of intersection.




Figure 5.7 On the left two subarcs of Γ intersect β with the same sign; on the right two subarcs of Γ intersect β with different signs.

Now, the support of the push-map Γ is a regular neighborhood of Γ in D_n . Hence we can decompose $(\beta)\Gamma$ as the union of subarcs of β that are fixed by Γ , which we refer to as **original subarcs**, together with **new subarcs** that lie in the support of Γ . Each of the new subarcs is one of four types: (I) joins an original subarc to a new point of intersection, (II) joins two new points of intersection forming part of the boundary of a disk bounding p_n , (III) joins two original subarcs, or (IV) joins two new points of intersection in a disk not containing only p_n .

Suppose now that the vertices s_1, s_2 of our bigon are both original points of intersection. Then the subarc of $(\beta)\Gamma$ that joins s_1 to s_2 cannot contain any new subarcs of type (I), type (II), or type (III) and must therefore contain a new subarc of type (IV). It follows that we have two subarcs $\beta' \subset \beta$ and $\alpha' \subset \alpha$ satisfying the following three statements:




Figure 5.8 Subarcs of types (I)-(IV): here original subarcs of β appear in red, while new subarcs of β are colored purple.

- the union $\alpha' \cup \beta'$ does not form a bigon;

- the union $\alpha' \cup (\beta')\Gamma$ forms a bigon; and
- the intersection $\alpha' \cap (\beta')\Gamma = \alpha' \cap \beta' = \{s_1, s_2\}$.

The disk bounded by $\alpha' \cup \beta'$ in D_n must contain at least one marked point, since we are assuming that α and β are in minimal position. If the disk bounded by $\alpha' \cup \beta'$ contained p_i for some $i \neq n$, then the disk bounded by $\alpha' \cup (\beta')\Gamma$ would also contain p_i , but this is not possible since $\alpha' \cup (\beta')\Gamma$ bounds a bigon. Therefore the disk bounded by $\alpha' \cup \beta'$ contains only the marked point p_n . We must also have that the final component Γ_F of Γ intersects β' in a single point, since $\alpha' \cup (\beta')\Gamma$ does not contain p_n . Moreover, no other subarc of Γ can intersect β' , as otherwise $\alpha' \cup (\beta')\Gamma$ would not form a bigon. In particular, it follows that Γ_I intersects $\alpha' \cup (\beta')\Gamma$ in a single point in α' . Hence the loop Γ must intersect the disk bounded by $\alpha' \cup \beta'$ as shown in Figure 5.9. However, such a configuration contradicts the fact that $\Phi \cdot \Gamma$ is a proper product. Therefore the vertices s_1 and s_2 cannot both be original points of intersection.




Figure 5.9 A new bigon involving two original points of intersection s_1 and s_2 between α and β must arise from a configuration of curves as depicted.

It remains to consider the case in which precisely one vertex of our bigon between α and $(\beta)\Gamma$ is an original point of intersection; we denote this vertex by s_1 . The other vertex s_2 must therefore be a new point of intersection formed when Γ meets a subarc of β and pushes it along until it crosses α . Without loss of generality, we can assume that s_2 is the new point of intersection to the left of the loop Γ , as shown in Figure 5.10.




Figure 5.10 The simplest case where a bigon between α and $(\beta)\Gamma$ involves precisely one original point of intersection.

Then there are two possible candidate regions for a bigon with s_2 as a vertex: one to the left of α as shown in the figure, and one to the right. In each case, the vertex s_1 lies along α on the same side of Γ as s_2 in the local picture shown in Figure 5.10. In the first case, there must be a piecewise Γ -arc forming a generalized β -to- α trigon with s_2 corresponding to its final point of intersection with α , which contradicts that $\Phi \cdot \Gamma$ is a proper product; see Figure 5.11. The second case also cannot occur, since this would imply that $(\beta)\Gamma$ continues on along a new subarc from s_2 to s_1 to form our bigon, which in turn implies that Γ itself must form a generalized bigon with β and that s_1 is also a new point of intersection, both of which are contradictions. The proposition follows.




Figure 5.11 A possible bigon that arises from a generalized trigon. The oriented arcs in gold are subarcs of a loop Γ .

In the next section, we will restrict our attention to the case $n = 4$. ■

6 Disk sequences in the 4-strand braid group

The aim of this section is to prove [Theorem 6.9](#) below, which states that the Burau representation ρ_4 is faithful on its restriction to the point-pushing group K_4 . By [Proposition 1.3](#), this will imply our [Main Theorem](#). Our strategy for proving [Theorem 6.9](#) is to use proper products to identify certain braids that satisfy the parity condition. This will allow us to apply Moody's theorem ([Theorem 2.3](#)) to establish faithfulness.

6.1 Disks and parity in B_4

We begin with a careful analysis of the types of disks that can arise in the disk sequence of a 4-strand braid. To that end, we introduce some useful notation: for a disk Δ contained in D_4 , we let $\mathcal{P}(\Delta)$ denote the set of marked points that are contained in the interior of Δ , and hence $\mathcal{P}(\Delta) \subseteq \{p_1, p_2, p_3, p_4\}$. We note that up to this point in the paper we have not used the assumption that $n = 4$. We will use that assumption in this section, particularly in [Lemmas 6.1](#) and [6.3](#).

Lemma 6.1 *Let Δ be a disk arising in the disk sequence of an element of B_4 . Then Δ is sign-preserving if and only if $\mathcal{P}(\Delta)$ is one of the following:*

$$\{p_1, p_3\}, \quad \{p_1, p_4\}, \quad \{p_2, p_3\}, \quad \{p_2, p_4\}, \quad \{p_1, p_3, p_4\}, \quad \{p_2, p_3, p_4\}.$$

Moreover, Δ is sign-changing if and only if $\mathcal{P}(\Delta)$ is one of the following:

$$\{p_3\}, \quad \{p_4\}, \quad \{p_3, p_4\}, \quad \{p_1, p_2, p_3\}, \quad \{p_1, p_2, p_4\}, \quad \{p_1, p_2, p_3, p_4\}.$$

Proof. Let Δ be a disk in the disk sequence of an element $\Phi \in B_4$, and let q and q' denote the two points of intersection between α and $\beta = (\beta_*^3)\Phi$ that lie in $\partial\Delta$. Without loss of generality, we assume that q lies between p_1 and q' on α , and that q precedes q' as we traveling along β .

Suppose that $\mathcal{P}(\Delta) = \{p_1\}$. Since β and α intersect transversely at q , the subarc of β joining the basepoint p_* to q must intersect $\partial\Delta$. Hence this subarc either intersects the β -component of $\partial\Delta$ or the subarc of α with endpoints q and q' . The former contradicts our assumption that β is simple, and the latter

contradicts our assumption that β and α are in minimal position. Therefore $\mathcal{P}(\Delta) \neq \{p_1\}$. By a similar argument, we see that $\mathcal{P}(\Delta) \neq \{p_2\}$ and that $\mathcal{P}(\Delta) \neq \{p_1, p_2\}$.

We next claim that Δ is sign-preserving if and only if it contains precisely one of the two endpoints of α ; the lemma follows immediately from the claim. Suppose first that Δ is sign-preserving. Without loss of generality, we assume that β is oriented down (in our standard picture where α is horizontal) at both q and q' , and that Δ contains p_1 ; see Figure 6.1.




Figure 6.1 A sign-preserving disk

Then $\partial\Delta$ is the union of a subarc of β and a subarc of α , both with endpoints q, q' , and the interior of Δ lies to the right of β as we travel along β from q to q' . This implies that Δ does not contain the subarc of α with endpoints q' and p_2 ; in particular Δ does not contain p_2 . A similar argument shows that if Δ is sign-changing and contains p_1 , then it also contains p_2 . The claim follows, and we are done. ■

6.2 The point-pushing subgroup K_4

Our focus in the remainder of this section will be the point-pushing subgroup K_4 . Our next step will be to analyze the types of disks that can appear in the disk sequence of certain braids in K_4 . Specifically, we will be constructing proper products of certain elements in K_4 to which we can apply Lemma 4.2. We will then analyze the corresponding disk sequences, which will ultimately provide a mechanism for applying Theorem 2.3.

Lemma 6.2 *Let $\Phi \in K_4$, let $\Gamma \in K_4$ be a push-map along a simple loop, and suppose that $\Phi = \Psi \cdot \Gamma$ is a proper product. Let $\Delta_1, \dots, \Delta_{m-1}$ denote the disk sequence of Φ . Then, for any fixed i , either Δ_i is a 1-disk containing only p_4 , or there is an isotopy of $\beta = (\beta_*^3)\Phi$ so that the β -component of the boundary Δ_i is disjoint from Γ .*

Before giving the proof, we emphasize that Lemma 6.2 only guarantees that the β -component of any single disk Δ_i can be isotoped off Γ . It does not guarantee that each such β -component can be isotoped off Γ simultaneously; indeed, this is not possible in general. We also introduce some notation that will be convenient in the proof of the lemma and in what follows: if γ is an arc or a loop in the disk D_n , we will refer to a disk whose boundary is the union of a single subarc of α together with a single subarc of γ as an α - γ **disk**.

Proof. (Lemma 6.2) Let $\beta' = (\beta_*^3)\Psi$. Again, we choose fixed representatives of Γ and α that are in minimal position. By Proposition 5.1, we may choose a representative of β' so that $(\beta')\Gamma$ is in minimal position with respect to α . Since Γ is a simple loop, we can represent the action of the push-map Γ via the schematic in Figure 6.2. Figure 6.2(b) shows the resulting local picture for $\beta = (\beta_*^3)\Phi$ after pushing β' along Γ , possibly creating a number of new disks in the associated disk sequence. For each such disk that does not contain p_4 , Figure 6.2(b) shows that its β -component is disjoint from Γ . For the remaining α - β disks in Figure 6.2(b), there is an isotopy of β in D_n that preserves the

disk sequence while removing the intersection of the β -component of these disks from Γ . This isotopy can be visualized by sliding the “topmost” part of any β -component of the boundary of a disk that intersects Γ in [Figure 6.2\(b\)](#) upwards along the rightmost subarc of α that appears in the local picture; one can perform a single isotopy that moves all such β -subarcs simultaneously off Γ .




Figure 6.2 Here blue arcs represent subarcs of α , while red/pink arcs represent subarcs of β . ■

A particular push-map. Let Γ_1 be the push-map corresponding to the loop shown in [Figure 6.3](#).




Figure 6.3 A particular push map Γ_1 .

Lemma 6.3 *Let $\Phi \in K_4$, let $\Gamma_1 \in K_4$ be the push-map along the simple loop shown in [Figure 6.3](#), and suppose that $\Phi = \Psi \cdot \Gamma_1$ is a proper product. Let Δ be a disk arising in the disk sequence of Φ . Then Δ is sign-preserving if and only if $\mathcal{P}(\Delta)$ is the following:*

$$\{p_2, p_3\}.$$

Moreover, Δ is sign-changing if and only if $\mathcal{P}(\Delta)$ is one of the following:

$$\{p_3\}, \quad \{p_4\}, \quad \{p_1, p_2, p_3\}, \quad \{p_1, p_2, p_3, p_4\}.$$

Proof. By [Lemma 6.2](#) any disk Δ in the disk sequence of Φ can have its β -component isotoped off of Γ_1 . Referring to [Figure 6.4](#), we see that, up to equivalence, there are only finitely many possible β -subarcs forming an α - β disk that is disjoint from Γ_1 . ■




Figure 6.4 The loop Γ_1 is shown in gold, together with red arcs indicating all possible disks arising in the disk sequence of a braid in K_4 such that the β -component of the disk can be isotoped off of Γ_1 .

The five possible types of disk that can arise up to equivalence are shown in [Figure 6.5](#) using our standard embedding of α and Γ_1 in D_4 .




Figure 6.5 The loop Γ_1 is shown in gold, together with five red arcs $\delta_1, \dots, \delta_5$ indicating all five types of disks arising in the disk sequence of a braid in K_4 such that the β -component of the disk can be isotoped off of Γ_1 .

6.3 Embedding in B_5

Under the hypotheses of [Lemma 6.3](#), there is only one possible type of disk arising in the disk sequence of a braid in K_4 that violates the parity condition: a disk containing all four punctures. In order to deal with this, we will pass briefly to the setting of the disk with five punctures. To that end, we now fix an embedding of D_4 in the disk D_5 ; we can think of this as simply fixing an additional marked point p_5 in the interior of the disk D_4 . If $\Phi \in B_4$, we let $f(\Phi) \in B_5$ denote the image of Φ in B_5 under the induced injection $f : B_4 \hookrightarrow B_5$.

The main goal of [Subsection 6.3](#) is to prove [Proposition 6.7](#) and its corollary, which (roughly speaking) establish conditions under which a push-map will satisfy the parity condition. We will begin with a series of lemmas that will equip us to deal with 4-disks, which under our hypotheses are the one remaining obstruction to satisfying the parity condition.

Let β be an arc in D_5 from p_* to p_3 , let Γ be a loop based at p_5 , and let Γ_F denote the subarc of Γ traveling from its final point of intersection with α back to p_5 . We note that the set of 4-disks in the disk sequence of Φ are necessarily nested; in particular, their respective β -components form a “band”, as indicated in [Figure 6.6](#). We will say the loop Γ is **disjoint from β modulo 4-disks** if Γ_F intersects this band of β -components transversely and Γ is otherwise disjoint from β , so that Γ intersects the β -component of each 4-disk precisely once and Γ is disjoint from the β -component of every other disk in the disk sequence of β .

Lemma 6.4 *Let $\Phi \in B_4$, let $\beta = (\beta_*^3)\Phi$, and let Γ be a loop in D_5 based at p_5 that is disjoint from β modulo 4-disks. Then the disk sequence of $f(\Phi) \cdot \Gamma$ is obtained from the disk sequence of Φ by replacing each 4-disk with a 5-disk. Moreover, each such 5-disk is necessarily sign-changing.*




Figure 6.6 The loop Γ is disjoint from β except for intersecting the “band” of β -components of the nested 4-disks. After performing the push-map Γ , the band of 4-disks is replaced with a band of 5-disks.

Before giving the proof, we emphasize that the loop Γ in the statement of [Lemma 6.4](#) need not be simple.

Proof. ([Lemma 6.4](#)) Since the only points of intersection between Γ and β lie in Γ_F , the effect on β of pushing along Γ is to push p_5 “into” each 4-disk, as shown in [Figure 6.6](#), while any other disks remain unchanged. As we have already observed, each 4-disk in the disk sequence of Φ is necessarily sign-changing, and hence the newly created 5-disks must also be sign-changing. ■

As above, we consider a loop Γ in D_5 based at p_5 , and let Γ_F denote the subarc of Γ traveling from its final point of intersection with α back to p_5 . Also, we refer to a disk whose boundary is the union of a single subarc of α together with a single subarc of Γ , where Γ is a loop, as an α - Γ **disk**. Finally, we say that Γ is **simple modulo 4-disks** if any self-intersection points of Γ occur as a point of intersection of Γ_F with the boundary of an α - Γ 4-disk.

Recall that Γ_I denotes the subarc of Γ from p_5 to its first point of intersection with α . If Γ is simple modulo 4-disks, we will say it is **properly simple modulo 4-disk** if the following three conditions hold: (1) each α - Γ disk is equivalent to one of the five disk types shown in [Figure 6.5](#), (2) there is some α - Γ disk bounding a 3-disk containing p_1, p_2 and p_3 , and (3) the arc Γ_I does not intersect β_*^3 .

Lemma 6.5 *Let Γ be a loop in D_5 based at p_5 that is properly simple modulo 4-disks, let $q_1, \dots, q_m$ denote the points of intersection of Γ with β_*^3 in order as we travel along β_*^3 , and let $\text{sign}(q_i)$ denote the sign of the intersection of Γ with β_*^3 at q_i . Then for all $1 \leq i < m$, we have that either $\text{sign}(q_i) = -\text{sign}(q_{i+1})$, or else the corresponding disks Δ_i and Δ_{i+1} are both 4-disks.*

In other words, [Lemma 6.5](#) says that the sign of the intersection of Γ with β_*^3 is alternating as we travel along β_*^3 , unless the intersections correspond to two adjacent 4-disks; for convenience we refer to this property as being **alternating modulo 4-disks**.

Proof. ([Lemma 6.5](#)) Since Γ is properly simple modulo 4-disks, it follows that p_5 is not contained in any α - Γ disks. Hence we can view p_5 as a basepoint on the boundary of the disk.

Let q_i and q_{i+1} be two points in $\Gamma \cap \beta_*^3$ that are adjacent along β_*^3 , and suppose that $\text{sign}(q_i) = \text{sign}(q_{i+1})$. Let β_i denote the subarc of β_*^3 from q_i to q_{i+1} , and similarly let γ_i denote corresponding subarc of Γ . Without loss of generality, we may assume that $\beta_i \cup \gamma_i$ is a circle containing at most 3 punctures, contained in the embedded copy of $D_4 \subset D_5$.

We now choose a point x that lies along Γ and between q_{i+1} and p_5 if Γ is simple, or otherwise between q_{i+1} and the first point of self-intersection of Γ as we travel along Γ ; see [Figure 6.7](#). The point x necessarily lies in the interior

of $\beta_i \cup \gamma_i$.




Figure 6.7 The disk bounded by $\beta_i \cup \gamma_i$.

Now, let γ_x denote the subarc of Γ from x to p_5 . Then γ_x necessarily intersects $\beta_i \cup \gamma_i$ in at least one point, which we denote by y . If $y \in \beta_i$ this contradicts our assumption that q_i and q_{i+1} are adjacent along β_*^3 . If $y \in \gamma_i$, then this contradicts our assumption that Γ is properly simple modulo 4-disks. This completes the proof that Γ is alternating module 4-disks with respect to β_*^3 . $\blacksquare$

The next lemma will be crucial for identifying braids that satisfy the parity condition.

Lemma 6.6 *Let Γ be a loop in D_5 based at p_5 that is properly simple modulo 4-disks. Then Γ satisfies the parity condition.*

Proof. Consider a disk Δ in the disk sequence of Γ . In what follows we will use the terminology of **original** and **new** subarcs and vertices, as introduced in the proof of [Proposition 5.1](#).

First, we note that both vertices of the disk Δ are new. Setting $\beta = (\beta_3^*)\Gamma$, we have moreover that the β -component of Δ is either a single new subarc, or it is the union of two new subarcs with an original subarc; see Figure 6.8.




Figure 6.8 The case where Δ is formed using two new subarcs and one original subarc.

In the former case, there are three possibilities: the disk Δ corresponds to an original α - Γ disk that is not a 4-disk, Δ is a 1-disk containing p_5 , or, by applying [Lemma 6.4](#), Δ is a sign-changing 5-disk. For each of these, Δ satisfies the parity condition.

By part (3) of the definition of being properly simple modulo 4-disks, β is the union of two new subarcs with an original subarc. Let γ_1 and γ_2 denote the two subarcs of Γ that correspond to the two new subarcs of β , where γ_1 precedes γ_2 in Γ . Suppose first that γ_1 and γ_2 both occur as subarcs of an α - Γ disks. If γ_1 and γ_2 are subarcs of equivalent α - Γ disks, then Δ corresponds to an original α - Γ -disk that is not a 4-disk, and hence Δ satisfies the parity condition.

Suppose next that γ_1 and γ_2 are subarcs of inequivalent α - Γ disks. We

begin by observing that both γ_1 and γ_2 are necessarily oriented from β_*^3 to α . If neither γ_1 nor γ_2 occur as subarcs of a 4-disk, then Δ does not contain p_5 . Moreover, Δ is equivalent to one of the disks shown in Figure 6.5 or is another sign-preserving 2-disk and hence satisfies the parity condition.

Referring again to Figure 6.8, we note that if γ_1 occurs as a subarc of a 4-disk, then γ_2 must occur as part of an α - Γ disk that is adjacent to the 4-disk as we travel along β_*^3 .

Since there is a 3-disk containing p_1, p_2, p_3 occurring as an α - Γ disk, and since no self-intersections of Γ can occur in this 3-disk, we see that γ_2 occurs either part of such a 3-disk or as part of a 1-disk containing p_4 ; see, respectively, the arcs δ_3 and δ_2 in Figure 6.5. By Lemma 6.5, the possible configurations that can occur are shown in Figure 6.9. Hence Δ is one of the following: a 1-disk containing p_4 , a 3-disk containing p_1, p_2 , and p_3 , or a sign-changing 5-disk; each of these satisfy the parity condition, and the proof is complete.




Figure 6.9 The four possible configurations: in each case γ_1 and γ_2 are highlighted in red, joined by an original subarc of β .

We emphasize that throughout what follows, the notation Γ_1 always refers to the particular push-map shown in Figure 6.3, while Γ refers to a push-map that generally depends on some other element $\Phi \in K_4$.

Proposition 6.7 *Let $\Phi \in B_4$, and suppose that Φ can be written as a proper product $\Phi = \Phi' \cdot \Gamma_1$. Then there exists a push-map $\Gamma \in K_5$ such that $f(\Phi) \cdot \Gamma$ is a proper product in B_5 and such that both $f(\Phi) \cdot \Gamma$ and Γ satisfy the parity condition. Furthermore we may choose Γ so that $\iota(\beta_*^3 \cdot (\Phi \cdot \Gamma), \alpha) \neq \iota(\beta_*^3 \cdot \Gamma, \alpha)$.*

Before we begin the proof of Proposition 6.7, we recall that an α - γ disk is a disk whose boundary is the union of a single subarc of α together with a single subarc of γ , where γ is a loop or an arc. In the case where γ is a union of two or more subarcs of a larger loop or arc Γ , we will also refer to such a disk as an α - Γ disk for convenience.

Proof. (Proposition 6.7) Suppose that the braid $\Phi \in B_4$ does not satisfy the parity condition. By Lemma 6.3 there must be at least one disk Δ in its disk sequence with $\mathcal{P}(\Delta) = \{p_1, p_2, p_3, p_4\}$. Moreover, any 4-disks that occur in its disk sequence are necessarily nested in D_4 . Consider the image $f(\Phi)$ of our braid in B_5 . We will construct a push-map $\Gamma \in K_5$ whose effect will be to replace each of the 4-disks in the disk sequence of $f(\Phi)$ with a 5-disk while preserving every other disk.

Let β denote the arc $(\beta_*^3)f(\Phi)$ in the disk D_5 , and choose a point q in the interior of the innermost 4-disk of β such that q is not contained in a k -disk in the disk sequence of $f(\Phi)$ for any $k \leq 3$. Since the arc β is simple, its

complement in the disk D_5 is connected, and hence we can choose a simple path γ_1 from p_5 to q such that $\iota(\gamma_1, \beta) = 0$; see Figure 6.10 for an example. We note that γ_1 necessarily intersects α , since q lies in the interior of the disk Δ and p_5 does not.




Figure 6.10 An example of the construction of the loop $\Gamma = \gamma_1 \cup \gamma_2$.

Next we choose a simple path γ_2 from q to p_5 that is disjoint from α and not isotopic to γ_1 , such that the intersection of γ_2 and β consists of precisely one point in the β -component of each 4-disk in the disk sequence of Φ . In other words, γ_2 travels transversely through the β -component of each 4-disk and then passes out of D_4 to p_5 without intersecting β again, and in particular without intersecting any other disks in the disk sequence of Φ ; again, see Figure 6.10 for an example. Finally, we define the loop $\Gamma \in K_5$ to be the product of paths $\gamma_2 \circ \gamma_1$.

By Lemma 6.4, the disk sequence of $(\beta)\Gamma = (\beta_*^3)(\Phi \cdot \Gamma)$ is exactly the same as that of $\beta = (\beta_*^3)\Phi$, except that each 4-disk in the original disk sequence of $\Phi \in B_4$ has been replaced by a 5-disk that is necessarily sign-changing. Hence $f(\Phi) \cdot \Gamma \in B_5$ satisfies the parity condition.

Consider next the disk sequence of $(\beta_*^3) \cdot \Gamma$. Note that Γ is properly simple modulo 4-disks; condition (1) is immediately satisfied by the construction, and condition (2) is satisfied as well from the fact that any proper product Φ yields a disk sequence containing a 3-disk bounding p_1, p_2, p_3 . By Lemma 6.6, this Γ satisfies the parity condition as well.

It remains to show that we can always choose Γ to ensure the following statement holds:

$$\iota(\alpha, (\beta_*^3) \cdot \Gamma) \neq \iota(\alpha, (\beta_*^3)(\Phi \cdot \Gamma)) \quad (2)$$

Suppose now that in our construction of Γ as above, we have $\iota(\alpha, (\beta_*^3) \cdot \Gamma) = \iota(\alpha, (\beta_*^3)(\Phi \cdot \Gamma))$, or in other words, $\iota(\alpha, (\beta_*^3)\Gamma) = \iota(\alpha, (\beta)\Gamma)$. Our strategy is to alter our construction of the loop Γ slightly so that the left-hand side of this equality increases, while the right-hand side remains unchanged.

To this end, we will form a new loop Γ' starting with the same choice of γ_1 as before, but we will choose a different arc γ'_2 joining the point q to p_5 to replace γ_2 , and set $\Gamma' = \gamma_1 \cup \gamma'_2$. We note that it is sufficient to find such an arc γ'_2 satisfying the following two properties:

1. as we traverse γ'_2 from q to p_5 , its first point of intersection with $\beta = (\beta_*^3) \cdot (\Phi)$ does not occur until after its final point of intersection with α ; and
2. $\iota(\alpha, (\beta_*^3) \cdot \Gamma') > \iota(\alpha, (\beta_*^3) \cdot \Gamma)$.

This follows from the fact that the first condition implies that $\iota(\alpha, (\beta)\Gamma') = \iota(\alpha, (\beta)\Gamma)$.

Revisiting the construction of Γ , let s denote the final point of intersection of γ_1 with α as we travel along γ_1 , and let $(\gamma_1)_F$ denote the subarc of γ_1 joining s to the point q . Since $D_4 \setminus (\beta \cup \gamma_1)$ is connected, there exists an arc γ' joining q to a point $r \in \alpha$, where the point r depends on the choice of γ' , such that γ' is disjoint from β and such that $\gamma' \cap \gamma_1 = \{q\}$.

We will now consider two cases. First, suppose that γ' can be chosen so that, the arc $\gamma' \cup (\gamma_1)_F$, together with the subarc of α joining r to s , forms a disk Δ containing p_1, p_2, p_3 , and p_4 .

We next look for an arc γ'' from r back to q satisfying the following:

1. the arc γ'' is disjoint from β ;
2. $\gamma'' \cap (\gamma_1 \cup \gamma') = \{q, r\}$; and
3. γ'' is not isotopic to γ' relative to the endpoints $\{q, r\}$.

This is equivalent to finding a nontrivial loop based at r in $D_4 \setminus (\beta \cup \gamma_1 \cup \gamma')$. This is always possible since $\beta \cap (\gamma_1 \cup \gamma') = \emptyset$, and hence $D_4 \setminus (\beta \cup \gamma_1 \cup \gamma')$ is homeomorphic to a disk with three punctures: p_1, p_2 , and p_4 . Finally, we set $\gamma'_2 = \gamma' \cup \gamma'' \cup \gamma_2$, and set $\Gamma' = \gamma_1 \cup \gamma'_2$. Since the subarc γ'_2 satisfies the two properties above, we have established the result in this case.

It remains to deal with the case where the arc γ' joining q to a point $r \in \alpha$ cannot be chosen so that $\gamma' \cup (\gamma_1)_F$ forms a 4-disk, and hence forms a disk containing a proper subset of $\{p_1, p_2, p_3, p_4\}$; recall that we are still using in this case that γ' is disjoint from β and that $\gamma' \cap \gamma_1 = \{q\}$.




Figure 6.11 An alternate construction of γ_2 : Case 1.

We will choose the arc γ' to travel from q to a point $r \in \alpha$ in parallel to the β -component of the boundary of the innermost 4-disk; there are two possibilities here. One choice will result in $(\gamma_1)_F \cup \gamma'$ forming an α - Γ disk containing k punctures, where $k < 4$, and the other choice for γ' results in an α - Γ disk containing $n - k$ punctures. Referring to [Figure 6.12](#), we see that here γ' is chosen to follow the innermost such β -component in the right-hand direction (as shown in the figure), so that $(\gamma_1)_F \cup \gamma'$ forms a 1-disk with α ; following to the left would form a 3-disk instead. In other words, we can assume without loss of generality that the α - Γ disk formed by $(\gamma_1)_F \cup \gamma'$ contains at most two punctures.




Figure 6.12 The horizontal lines in this figure represent β -components of 4-disks. The choice of the arc γ' shown here yields a disk with 1 puncture, by traveling to the right in parallel with the β -component of the innermost 4-disk.

Our construction now reduces as in the previous case to finding a nontrivial loop based at r in $D_n \setminus (\beta \cup \gamma_1 \cup \gamma' \cup \gamma'')$. Again, this is always possible since $\beta \cap (\gamma_1 \cup \gamma') = \emptyset$, and since by construction $D_4 \setminus (\beta \cup \gamma_1 \cup \gamma')$ is homeomorphic to a disk with at least one puncture. This completes the proof of the proposition. $\blacksquare$

We remark that the same construction used to find the loop Γ in the proof of Proposition 6.7 could be iterated to find, for any n , and for any $\Phi \in B_n$, some $\Gamma_1, \dots, \Gamma_k$ so that, for the appropriate inclusion map f , we have that $f(\Phi) \cdot \Gamma_1 \dots \Gamma_k \in B_{n+k}$ satisfies the parity condition. The rough idea here is to start with the innermost disk of each occurring type that does not satisfy the parity condition, and find a suitable push-map that “corrects” its parity. However, in this case we cannot guarantee that the braid $\Gamma_1 \dots \Gamma_k$ satisfies the parity condition. In our case, the fact the braid Γ also satisfies the parity condition is important in the proof of our Main Theorem. The fact that we only need to correct one type of disk, so we only require one push map Γ , is crucial. In the general case where $k > 1$, we would not be able to simultaneously assume that Γ_2 does not cross both β and $(\beta) \cdot \Gamma_1$, and so this argument would fail.

The following corollary of Proposition 6.7 enables us to derive conditions on the Moody polynomial of a push-map $\Phi \in K_4$, as part of a proper product.

Corollary 6.8 *If $\Phi \in K_4$ admits a proper product $\Phi = \Phi' \cdot \Gamma_1$ and satisfies $\iota(\alpha, (\beta_*^3)\Phi) > 0$, then there exists a push-map $\Gamma \in K_5$ such that $\mathbb{M}_{f(\Phi) \cdot \Gamma} \neq \mathbb{M}_\Gamma$.*

Proof. Let Φ be a braid in K_4 , and let Γ be an associated push-map in K_5 as in the statement of Proposition 6.7. By Proposition 6.7, both Γ and $\Phi \cdot \Gamma$ satisfy the parity condition. It then follows from Lemma 4.2 that there are no cancellations in either of the corresponding Moody polynomials $\mathbb{M}_\Gamma$ and $\mathbb{M}_{f(\Phi) \cdot \Gamma}$.

Let M_- and M_+ be the lowest and highest exponents that appear in the Moody polynomial $\mathbb{M}_{f(\Phi) \cdot \Gamma}$, and similarly let M'_- and M'_+ denote the corresponding values for $\mathbb{M}_\Gamma$. Then, since neither polynomial has any cancellations, we can write the corresponding Moody polynomials as follows:

$$\begin{aligned} \mathbb{M}_{f(\Phi) \cdot \Gamma} &= \sum_{j=M_-}^{M_+} a_j t^j = |a_{M_-}| + \dots + |a_{M_+}| = \iota(\alpha, (\beta_*^3)(f(\Phi) \cdot \Gamma)); \quad \text{and} \\ \mathbb{M}_\Gamma &= \sum_{j=M'_-}^{M'_+} b_j t^j = |b_{M'_-}| + \dots + |b_{M'_+}| = \iota(\alpha, (\beta_*^3)\Gamma). \end{aligned}$$

By Proposition 6.7 we can assume that $\iota(\alpha, (\beta_*^3)(f(\Phi) \cdot \Gamma)) \neq \iota(\alpha, (\beta_*^3)\Gamma)$. The result follows. $\blacksquare$

6.4 Faithfulness on Brun_4

As described in [Section 1](#), the proof of our [Main Theorem](#) follows from the next result.

Theorem 6.9 *The Burau representation ρ_4 is faithful on its restriction to Brun_4 .*

Proof. Let Φ denote a nontrivial element in Brun_4 . We make the following elementary observation: $\Phi \in \text{Brun}_4$ lies in the kernel of the Burau representation ρ_4 if and only if, for every element $y \in B_4$, the conjugate $y\Phi y^{-1}$ also lies in the kernel of ρ_4 . We also recall from our discussion in [Section 1](#) that any nontrivial element of Brun_4 is pseudo-Anosov, and that the image of the arc β_*^3 under a pseudo-Anosov map necessarily intersects the arc α nontrivially.

Given a simple loop such as Γ_1 the only possible bigon-forming polygons that can occur are β -to- α trigons and rectangles. Now, β -to- α trigons can be resolved by an isotopy of β , and rectangles are not possible in this case since the points of intersection between α and Γ_1 alternate in sign as we travel along α .

Furthermore, if we consider Φ and Γ_1 as freely reduced words in the free group K_4 with respect to the free basis shown in [Figure 6.13](#), and if the juxtaposition of the two words (yielding the product $\Phi \cdot \Gamma_1$) is also a freely reduced word with respect to this free basis, then the second condition for $\Phi \cdot \Gamma_1$ to be a proper product is immediately satisfied. One way to see this is to note that, as a freely reduced word in this basis, Γ_1 begins with y_2^{-1} . If $\Phi \cdot \Gamma_1$ is freely reduced, then Φ cannot end with y_2 ; if it does, we can replace Φ with an appropriate conjugate that does not end with y_2 . Thus we may assume that Φ has been conjugated so that it can be written as a proper product $\Phi = \Phi' \cdot \Gamma_1$ for some $\Phi' \in K_4$.




Figure 6.13 A free basis for K_4 .

By [Corollary 6.8](#), there is some push-map $\Gamma \in K_5 < B_5$ so that $\mathbb{M}_{f(\Phi) \cdot \Gamma} \neq \mathbb{M}_\Gamma$. It now follows from [Theorem 2.3](#) that $f(\Phi)$ does not lie in the kernel of ρ_5 , and hence Φ does not lie in the kernel of ρ_4 . Hence the Burau representation of B_4 is faithful on the point-pushing subgroup Brun_4 . ■

7 An example

We end with an illustrative example of a push-map $\Gamma \in K_5$ of the form guaranteed by [Proposition 6.7](#). We will apply Γ to an arc β whose corresponding Moody polynomial admits cancellation, and we will see that after applying Γ there is no longer any cancellation. Let β denote the arc shown in [Figure 7.1](#). We recall that an arc β does not uniquely determine a braid Φ such that $\beta = (\beta_*^3)\Phi$; nevertheless in what follows it will be useful to make a choice of such a Φ to simplify notation.

To simplify the calculation while explaining the important parts of the construction, we have chosen to specify an arc β whose disk sequence satisfies the same properties as those guaranteed by [Lemma 6.3](#), but our choice of β does not arise as $(\beta_*^3)\Phi$ for any choice of $\Phi \in K_4$. (Recall that we restricted to K_4 in [Theorem 6.9](#) in order to ensure that Φ can be written as a proper

product $\Phi = \Phi' \cdot \Gamma_1$, which allows us to apply [Lemma 6.3](#).) We make this choice for illustrative purposes because the simplest example of an arc β admitting cancellations, and arising from a proper product of the form $\Phi = \Phi' \cdot \Gamma_1 \in K_4$, would have many more intersections.




Figure 7.1 An example of an arc β whose Moody polynomial admits a cancellation.

Note that any braid Φ inducing the arc β does not satisfy the parity condition. Here and throughout this appendix, we write the terms of the unsimplified Moody polynomial in order corresponding to points of intersection of α with β as we travel along β from the basepoint to p_3 . We compute the Moody polynomial for Φ as follows, noting that the linear terms appear with opposite signs, yielding a cancellation.

$$\begin{aligned} \mathbb{M}_\Phi &= t^0 + t^2 + t^4 - t^3 - t + t^{-3} + t^{-1} + t \\ &= 1 + t^2 + t^4 - t^3 + t^{-3} + t^{-1}. \end{aligned}$$

Following the procedure given in the proof of [Proposition 6.7](#), we obtain the loop $\Gamma \in K_5$ shown in [Figure 7.2](#). Pushing β along Γ gives us $\beta \cdot \Gamma$ as shown in [Figure 7.3](#), which now satisfies the parity condition. We next compute the Moody polynomial of the arc $\Phi \cdot \Gamma$:

$$\begin{aligned} \mathbb{M}_{\Phi \cdot \Gamma}(t) &= t^0 + t^2 + t^4 - t^3 - t + t^{-4} + t^{-2} + t^0 \\ &= 2 + t^2 + t^4 - t^3 - t + t^{-4} + t^{-2}. \end{aligned}$$




Figure 7.2 The push map $\Gamma \in K_5$ shown in gold, with the point q marked from the proof of [Proposition 6.7](#).




Figure 7.3 The image $(\beta)\Gamma$ in D_5

Comparing with our calculation of the unsimplified Moody polynomial of Φ above, we see that the effect of Γ on the Moody polynomial has been to increase the exponent of each the last three terms (those following the appearance of the 5-disk) by 1, and there is no longer any cancellation.

Furthermore, the arc $(\beta_*^3) \cdot \Gamma$ determining the Moody polynomial of Γ is shown in Figure 7.4 and satisfies the parity condition as well. The unsimplified Moody polynomial for the arc $(\beta_*^3) \cdot \Gamma$ has 20 terms; the key point here is that any β -arc satisfying the parity condition may have any number of terms appearing with equal exponents, but every term with the same exponent will have the same sign.

$$\begin{aligned} \mathbb{M}_\Gamma(t) &= -t^0 - t^2 + t + t^{-1} - t^0 - t^2 + t^3 + t + t^{-1} - t^{-2} \\ &\quad - t^0 + t^{-1} + t^{-3} - t^{-2} - t^0 + t + t^{-1} + t^{-3} - t^{-2} - t^0 \\ &= 2t^{-3} - 3t^{-2} + 4t^{-1} - 5 + 3t - 2t^2 + t^3. \end{aligned}$$




Figure 7.4 The image $(\beta_*^3) \cdot \Gamma$ in D_5 .

References

- [1] Alperin, Roger C.; Benson Farb and Guennadi A. Noskov, *A Strong Schottky Lemma for Nonpositively Curved Singular Spaces*, Geometriae Dedicata **92**(1) (2002) 235-243, doi:10.1023/A:1019650112871.
- [2] Birman, Joan S. and Tara E. Brendle, *Braids: a survey*, Handbook of knot theory (2005) 19-103, Elsevier B. V., Amsterdam, doi:10.1016/B978-044451452-3/50003-4.
- [3] Bigelow, Stephen, *The Burau representation is not faithful for $n = 5$* , Geom. Topol. **3** (1999) 397-404, doi:10.2140/gt.1999.3.397.

- [4] Birman, Joan S., *Braids, links, and mapping class groups*, Annals of Mathematics Studies no. 82 (1974), Princeton University Press; University of Tokyo Press, Princeton, N.J.; Tokyo, doi:10.1515/9781400881420.
- [5] Beridze, A. and P. Traczyk, *Burau representation for $n = 4$* , J. Knot Theory Ramifications **27**(3) (2018) 1840002, 6, doi:10.1142/S0218216518400023.
- [6] Beridze, A. and P. Traczyk, *Forks, noodles and the Burau representation for $n = 4$* , Trans. A. Razmadze Math. Inst. **172**(3) (2018) 337-353, doi:10.1016/j.trmi.2018.05.001.
- [7] Burau, Werner, *Über Zopfgruppen und gleichsinnig verdrillte Verkettungen*, Abh. Math. Sem. Univ. Hamburg **11**(1) (1935) 179-186, doi:10.1007/BF02940722.
- [8] Calvez, Matthieu and Tetsuya Ito, *A Garside-theoretic analysis of the Burau representations*, J. Knot Theory Ramifications **26**(7) (2017) 1750040, 29, doi:10.1142/S0218216517500407.
- [9] Datta, Amitesh, *A strong characterization of the entries of the Burau matrices of 4-braids: The Burau representation of the braid group B_4 is faithful almost everywhere*, Preprint (2022).
- [10] Dlugie, Ethan, *The Burau representation and shapes of polyhedra*, Algebr. Geom. Topol. **24**(5) (2024) 2787-2805, doi:10.2140/agt.2024.24.2787.
- [11] Farb, Benson and Dan Margalit, *A primer on mapping class groups*, Princeton Mathematical Series **49** (2012), Princeton University Press, Princeton, NJ, doi:10.23943/princeton/9780691147949.001.0001.
- [12] Fullarton, Neil J. and Richard Shadrach, *Observed periodicity related to the four-strand Burau representation*, Exp. Math. **28**(1) (2019) 57-66, doi:10.1080/10586458.2017.1335625.
- [13] Gibson, Joel; Geordie Williamson and Oded Yacobi, *4-strand Burau is unfaithful modulo 5*, J. Exp. Math. **1**(1) (2025) 36-50, doi:10.56994/jxm.001.001.003.
- [14] Jones, V. F. R., *Hecke algebra representations of braid groups and link polynomials*, Annals of Mathematics **126**(2) (1987) 335-388, doi:10.2307/1971403.
- [15] Kasahara, Yasushi, *Remarks on the faithfulness of the Jones representations*, Groups of diffeomorphisms, Adv. Stud. Pure Math. **52** (2008) 369-381, Math. Soc. Japan, Tokyo, doi:10.2969/aspm/05210369.
- [16] Long, D. D., *A note on the normal subgroups of mapping class groups*, Math. Proc. Cambridge Philos. Soc. **99**(1) (1986) 79-87, doi:10.1017/S0305004100063957.
- [17] Long, D. D. and M. Paton, *The Burau representation is not faithful for $n \geq 6$* , Topology **32**(2) (1993) 439-447, doi:10.1016/0040-9383(93)90030-Y.
- [18] Lee, Sang Jin and Won Taek Song, *The kernel of $Burau(4) \otimes_{\mathbb{Z}_p}$ is all pseudo-Anosov*, Pacific J. Math. **219**(2) (2005) 303-310, doi:10.2140/pjm.2005.219.303.
- [19] Margalit, Dan, *Problems, questions, and conjectures about mapping class groups*, Breadth in contemporary topology, Proc. Sympos. Pure Math. **102** (2019) 157-186, Amer. Math. Soc., Providence, RI, doi:10.1090/pspum/102/12.
- [20] Moody, John Atwell, *The Burau representation of the braid group B_n is unfaithful for large n* , Bull. Amer. Math. Soc. (N.S.) **25**(2) (1991)

379-384, doi:10.1090/S0273-0979-1991-16080-5.

- [21] Moran, Siegfried, *Some free groups of matrices and the Burau representation of B_4* , Math. Proc. Cambridge Philos. Soc. **110**(2) (1991) 225-228, doi:10.1017/S0305004100070298.
- [22] Magnus, Wilhelm and Ada Peluso, *On a theorem of V. I. Arnold*, Comm. Pure Appl. Math. **22** (1969) 683-692, doi:10.1002/cpa.3160220508.
- [23] Whittlesey, Kim, *Normal all pseudo-Anosov subgroups of mapping class groups*, Geom. Topol. **4** (2000) 293-307, doi:10.2140/gt.2000.4.293.
- [24] Witzel, Stefan and Matthew C. B. Zaremsky, *A free subgroup in the image of the 4-strand Burau representation*, J. Knot Theory Ramifications **24**(12) (2015) 1550065, 16, doi:10.1142/S0218216515500650.