

An Introduction to Braid Theory

Maurice Chiodo

November 4, 2005

Abstract

In this paper we present an introduction to the theory of braids. We lay down some clear definitions of a braid, and proceed to establish the braid group \mathbf{B}_n . We show that both the word and conjugacy problems are solvable on this group, and conclude with theorems by Alexander and Markov to develop a correspondence between braids and links.

Acknowledgements

I would like to thank my supervisor Lawrence Reeves, whose guidance has proven extremely valuable over the course of this year. It has been a pleasure to work on such an interesting topic, alongside such an excellent mentor.

Contents

1	Introduction	5
2	Definitions	10
2.1	Braids	10
2.2	Braid equivalence	11
2.3	Braid diagrams	16
3	The braid group	20
3.1	Braids as a group	20
3.2	A presentation of the braid group	27
4	Properties of the braid group	37
4.1	Some results about the braid group	37
4.2	Braid invariants	42
4.3	Pure braids	44
4.4	Quotients of \mathbf{B}_n	46
5	The word and conjugacy problems on \mathbf{B}_n	50
5.1	The word problem	50
5.2	The conjugacy problem	59
6	Braids and links	62
6.1	Braid closure	62

6.2	Alexander's theorem	65
6.3	Markov's theorem	69
7	Conclusion	76
7.1	Summary	76
7.2	Further remarks	76
	References	78

1 Introduction

We begin by giving an informal introduction, intended to give the reader an intuitive understanding of braids.

The theory of braids has been studied since the early 1920's, founded by Emil Artin, a German mathematician (see [1], which is his original paper "*Theorie der Zöpfe*" from 1925). Though his studies were initially motivated by the geometric constructions of braids, it was not long before the powerful algebra behind braid theory became evident. Since then the theory has branched out into many fields of application, from encryption to solving polynomial equations. However, the study of braids in themselves is mathematically both rich and deep, being an extension of a concept that even a child can understand. We begin by giving an intuitive description of braids, to help motivate our definitions and ideas later on. So, what is a braid? Essentially, a braid is a geometric object that can, after some work, be viewed algebraically. We begin by taking a unit cube, and in it we place n strands of string, subject to the following conditions:

1. No part of any strand lies outside the cube.
2. Each strand begins on the top face of the cube, and ends on the bottom face.
3. No two strands intersect.
4. As we traverse any strand from the top face, we are always moving downwards. This means that no strand has any horizontal segment, or any segment that 'loops up'.

The resulting collection of strands is called an **n -braid**. Figure 1 shows some examples of 3-braids, while Figure 2 shows some non-examples of 3-braids. Now, given this loose definition of a braid, we can develop some sort of equivalence of braids. Given an n -braid β (in the unit cube) we say it is equivalent to another n -braid β' if the strands of β can be perturbed to the strands of β' without doing any of the following:

1. Moving any part of any strand out of the cube.
2. Cutting any strand.
3. Moving any endpoint of any strand.




Figure 1: *Two 3-braids.*




Figure 2: *Neither of these are braids, as they violate the conditions in the definition.*

So imagine the strands in a box, with their ends glued to the top/bottom of the box. We are, in essence, only allowed to ‘shake the box’. Figure 3 shows an example of such a move. It often becomes quite tedious and cumbersome to draw the cube around every braid, so for simplicity we omit it, and instead draw a projection of the braid onto the plane.




Figure 3: *Perturbing the strands of a braid.*

We can also multiply two n -braids β and β' by joining the bottom of β to the top of β' . By doing this we create a new n -braid which we shall denote by $\beta\beta'$. Figure 4 gives an example of this. It turns out that, for any given $n \in \mathbb{N}$, the set of equivalence classes of n -braids form a finitely-presented group, called the **n -braid group \mathbf{B}_n** , as we show in section 3. Moreover, both the word and conjugacy problems are solvable for this group, and we show this in section 5.




Figure 4: *Two braids, and their product.*

Any given n -braid β can be turned into an oriented link in a very intuitive manner. We simply draw a projection of the braid, and successively add arcs around the side of β to ‘close’ the open ends of the braid. Each arc connects the start and end of a strand, so we use n arcs in total. This is called the **closure** of the braid β , and is denoted by $\tilde{\beta}$. We give an example of a braid and its closure in Figure 5. The orientation on $\tilde{\beta}$ is given by orienting each strand of the original braid β with an arrow from top to bottom, then continuing these arrows around.

Figure 6 shows the closure of a braid with orientation drawn in. Conversely, a theorem by Alexander shows us that given any oriented link L , we can find some braid β such that L is the closure of β . However, this reverse procedure is much more difficult than finding the closure of a braid, as we shall see in section 6.2. We give an example of an oriented link drawn as the closure of a braid in Figure 7.




Figure 5: *A braid and its closure.*




Figure 6: *The induced orientation on the closure of a braid.*




Figure 7: *An oriented link, drawn again as the closure of a braid.*

2 Definitions

We now make our our intuitive ideas of a braid more precise, by introducing formal definitions. This enables us to convey our ideas with no ambiguity, and thus prove some interesting and insightful results.

2.1 Braids

We begin by giving a formal definition of a braid, as well as describing a convenient way to draw a braid. We are then in a position to go on to explore some direct consequences of our definition.

Definition 2.1.

Define a **level plane** $E_s \subseteq \mathbb{R}^3$ by

$$E_s := \{(x, y, z) \in \mathbb{R}^3 \mid z = s\}$$

Thus E_s is the infinite horizontal plane that intersects the z -axis at $z = s$.

Definition 2.2.

Let \mathbb{D} be the unit cube in the positive octant of Euclidean 3-space, with one vertex at the origin. So $\mathbb{D} = \{x, y, z \in \mathbb{R} : 0 \leq x, y, z \leq 1\}$. We define n points A_1, \dots, A_n on the top face of \mathbb{D} by

$$A_i := \left(\frac{1}{2}, \frac{i}{n+1}, 1 \right), \quad 1 \leq i \leq n$$

Similarly, we define n points B_1, \dots, B_n on the bottom face of \mathbb{D} by

$$B_i := \left(\frac{1}{2}, \frac{i}{n+1}, 0 \right), \quad 1 \leq i \leq n$$

Figure 8 illustrates our set up of \mathbb{D} . We now add n polygonal arcs d_1, \dots, d_n to \mathbb{D} such that the following hold:

1. The arcs d_1, \dots, d_n are mutually disjoint.
2. Each d_i begins at some A_j and ends at some B_k .
3. For any $0 \leq s \leq 1$ and any $1 \leq i \leq n$, $E_s \cap d_i$ is exactly one point.
4. Each d_i is contained entirely in \mathbb{D} .




Figure 8: *The unit cube \mathbb{D} with vertex at the origin.*




Figure 9: *Examples of 2, 3 and 4-braids.*

The resulting collection of n arcs d_1, \dots, d_n is called an **n -braid**. For any given $1 \leq i \leq n$, d_i is called the i^{th} **braid string** (though we may refer to them as strands or arcs). For a given $n \in \mathbb{N}$, the set of all n -braids is denoted \mathcal{B}_n .

To clarify, condition 3 of the above definition ensures that our braid strings are strictly decreasing. In Figure 9 we give an example of 2, 3 and 4-braids. Note that, for simplicity, we will draw our braid strings as smooth arcs, but we must remember that we are actually dealing with polygonal arcs here.

2.2 Braid equivalence

We now wish to give a concrete definition of what we mean by braid equivalence. Intuitively, two braids are the same if we can jiggle one to make it look like the other. With our new definition of a braid, we can clearly explain what we mean by this. So we begin by defining a set of fundamental moves that we can perform on a braid, that equates to our intuitive idea of jiggling.




Figure 10: *The elementary move Ω ; replacing the existing edge AB with the edges $AC \cup BC$.*

Definition 2.3.

Let $\beta \in \mathcal{B}_n$ be an n -braid, and let AB be an edge of some braid string d of β . Let C be another point in \mathbb{D} such that the solid triangle ABC has the following properties:

1. No other braid string of β meets ABC .
2. ABC meets d only along AB .
3. For any $0 \leq s \leq 1$, the level plane E_s meets $AC \cup BC$ in at most one point.

Then we shall define an operation Ω as follows: Replace the edge AB in d by the two edges $AC \cup BC$, as shown in Figure 10. If instead edges AC and BC are in d , and the level planes E_s each meet AB in at most one point, then we can also define an operation Ω^{-1} in the reverse manner: Replace the edges $AC \cup BC$ in d by AB , as shown in Figure 11. The moves Ω and Ω^{-1} are called **elementary moves** on the braid β .

Theorem 2.4.

If β is an n -braid and we perform an elementary move on β to obtain a collection of arcs β' , then β' is also an n -braid.

Proof.

Any braid string in β altered by an elementary move remains contained in \mathbb{D} (Since A, B, C all lie in \mathbb{D}), and condition 3 of Definition 2.3 ensures that elementary moves do not create extra intersections of strings with level planes. At no point do we remove a string entirely, nor do we introduce any new strings.




Figure 11: *The elementary move Ω^{-1} ; replacing the existing edges $AC \cup BC$ with the edge AB .*

And the start/end points of the strings in β remain fixed. Hence our resultant collection of strings β' is an n -braid. \square

Definition 2.5.

Let β, β' be n -braids. Suppose there exists a finite sequence of elementary moves that transform β to β' . Then β is said to be **equivalent** to β' , and is denoted by writing $\beta \sim \beta'$.

Note here that all elementary moves are, by definition, performed *inside* \mathbb{D} .

Theorem 2.6.

The relation \sim is an equivalence relation on \mathcal{B}_n .

Proof.

1. Reflexivity.

Let $\beta \in B_n$. Then the finite sequence of no elementary moves transforms β to β , and so $\beta \sim \beta$.

2. Symmetry.

Let $\beta, \beta' \in B_n$ such that $\beta \sim \beta'$. Then there exists a finite sequence of elementary moves $\Omega_1^{\epsilon_1}, \dots, \Omega_m^{\epsilon_m}$ that transforms β to β' . So we have $\beta \xrightarrow{\Omega_1^{\epsilon_1}} \dots \xrightarrow{\Omega_m^{\epsilon_m}} \beta'$. Reversing the sequence gives $\beta' \xrightarrow{\Omega_m^{-\epsilon_m}} \dots \xrightarrow{\Omega_1^{-\epsilon_1}} \beta$. Thus the finite sequence of elementary moves $\Omega_m^{-\epsilon_m}, \dots, \Omega_1^{-\epsilon_1}$ transforms β' to β (i.e., $\beta' \sim \beta$).

3. Transitivity.

Let $\beta, \beta', \beta'' \in B_n$ such that $\beta \sim \beta', \beta' \sim \beta''$. Then there exists a finite sequence

of elementary moves $\Omega_1^{\epsilon_1}, \dots, \Omega_m^{\epsilon_m}$ that transforms β to β' , and a finite sequence of elementary moves $\bar{\Omega}_1^{\bar{\epsilon}_1}, \dots, \bar{\Omega}_k^{\bar{\epsilon}_k}$ that transforms β' to β'' . So we have $\beta \xrightarrow{\Omega_1^{\epsilon_1}} \dots \xrightarrow{\Omega_m^{\epsilon_m}} \beta'$ and $\beta' \xrightarrow{\bar{\Omega}_1^{\bar{\epsilon}_1}} \dots \xrightarrow{\bar{\Omega}_k^{\bar{\epsilon}_k}} \beta''$. This gives the sequence $\beta \xrightarrow{\Omega_1^{\epsilon_1}} \dots \xrightarrow{\Omega_m^{\epsilon_m}} \beta' \xrightarrow{\bar{\Omega}_1^{\bar{\epsilon}_1}} \dots \xrightarrow{\bar{\Omega}_k^{\bar{\epsilon}_k}} \beta''$. Thus the finite sequence of elementary moves $\Omega_1^{\epsilon_1}, \dots, \Omega_m^{\epsilon_m}, \bar{\Omega}_1^{\bar{\epsilon}_1}, \dots, \bar{\Omega}_k^{\bar{\epsilon}_k}$ transforms β to β'' (i.e., $\beta \sim \beta''$). \square

We now introduce three other equivalence relations on braids.

Definition 2.7.

Let β, β' be n -braids in \mathbb{D} . Suppose there exists a homeomorphism $H : \mathbb{D} \times [0, 1] \rightarrow \mathbb{D} \times [0, 1]$ of the form $H(x, t) = (h_t(x), t)$ for $x \in \mathbb{D}$ and $0 \leq t \leq 1$. Also suppose that, for all $t \in [0, 1]$, the homeomorphism $h_t : \mathbb{D} \rightarrow \mathbb{D}$ satisfies the following:

1. $h_t|_{\partial\mathbb{D}} = id : \partial\mathbb{D} \rightarrow \partial\mathbb{D}$
2. $h_0 = id : \mathbb{D} \rightarrow \mathbb{D}$
3. $h_1(\beta) = \beta'$

Then β is said to be **ambient isotopic** to β' , denoted $\beta \approx \beta'$. Such a homeomorphism H is called an **ambient isotopy**.

The idea behind ambient isotopy is that, as well as deforming β to its homeomorphic image β' , we also preserve the structure of the surrounding space. Given some $n \in \mathbb{N}$, it is easy to show that any two n -braids are homeomorphic as topological subspaces of \mathbb{R}^3 . The way we distinguish them is to look at the structure of the remainder of \mathbb{D} .

Definition 2.8.

Let β, β' and H be as in Definition 2.7. Suppose that H also satisfies the following condition:

4. For all $t \in [0, 1]$, $h_t(\beta)$ is an n -braid.

Then β is said to be **strong isotopic** to β' , denoted by $\beta \sim_s \beta'$, and the homeomorphism H is called a **strong isotopy**, or just **s-isotopy**.

Strong isotopy is just a stronger form of ambient isotopy, where we impose the condition that the image of β must always be a braid throughout the deformation. We will see later that if two braids are ambient isotopic, then they are in fact strong isotopic. In other words, being able to deform one braid into another implies that there exists a ‘better way’ to do it, whereby at each step of the deformation we still have a braid.

Definition 2.9.

Let β, β' be n -braids in \mathbb{D} . Suppose there exists a homeomorphism $h : \mathbb{D} \rightarrow \mathbb{D}$ such that the following conditions hold:

1. $h(\beta) = \beta'$
2. $h|_{\partial\mathbb{D}} = id : \partial\mathbb{D} \rightarrow \partial\mathbb{D}$

Then β is said to be **h -equivalent** to β' , denoted $\beta \sim_h \beta'$, and the homeomorphism h is called an **h -equivalence**.

All we really mean by h -equivalence is that we send \mathbb{D} to itself via a homeomorphism such that β is sent to β' , and the boundary of \mathbb{D} is unchanged.

It turns out that all the relations we have defined so far tie in nicely with our idea of equivalence of braids, as the next theorem shows.

Theorem 2.10.

Ambient isotopy, strong isotopy and h -equivalence are all equivalence relations.

Proof.

The proof of this is simple topology, giving us little insight into the bigger picture of braid theory. We leave it as an exercise. See [4], pp 96-107. □

Intuition tells us that these definitions are all essentially saying the same thing. Namely, that two n -braids β, β' are the same if we can jiggle one in \mathbb{D} to make it look like the other. The above equivalence relations are merely formalised ways of saying we can jiggle braids, and though they may appear different, it turns out that they are not. That is, two braids equivalent under one of the above relations are equivalent under all others. The following theorem asserts this (though we omit the proof here and instead refer the reader to an appropriate reference).

Theorem 2.11.

Let β, β' be n -braids in \mathbb{D} . Then the following statements are equivalent:

1. β is equivalent to β' (i.e., $\beta \sim \beta'$).
2. β is ambient isotopic to β' (i.e., $\beta \approx \beta'$).
3. β is h -isotopic to β' (i.e., $\beta \sim_h \beta'$).
4. β is strong-isotopic to β' (i.e., $\beta \sim_s \beta'$).

Proof.

See [4], pp 96-107 for a full proof. □

2.3 Braid diagrams

Now that we have a more diverse (yet equivalent) set of equivalence relations on \mathcal{B}_n , we can be more relaxed in our interpretation of a braid. However, we have yet to formalise a simple visual presentation of a braid. Drawing \mathbb{D} each time we wish to describe a braid can become cumbersome and confusing, especially when we have a braid on a large number of strings, or with a large number of crossings. So we define a simple way to project a braid onto the plane, that preserves enough information to allow us to recover the braid (up to equivalence). In the process, we uncover a way to generate braids from simple building blocks, which we shall see later.

Definition 2.12.

Let $\beta \in B_n$ be in \mathbb{D} . We begin by retracting \mathbb{D} to the yz plane via the projection map $p : \mathbb{D} \rightarrow \mathbb{D}$, $p(x, y, z) = (0, y, z)$. The effect of this is to squash the braid strings d_1, \dots, d_n of β onto the back face of \mathbb{D} . For simplicity, we will treat our page as the yz plane, and denote $p(\beta)$ by $\hat{\beta}$. See Figure 12 for an example of a braid and its projection. The projection $p(\beta)$ gives a set of n curves $\hat{d}_1, \dots, \hat{d}_n$ in the plane, each \hat{d}_i being the image of d_i under p . These curves may have many (or possibly infinite) points of intersection, called **intersection points**. Any projection $\hat{\beta}$ of β satisfying the following three conditions is said to be a **regular projection** of β :

1. $p(\beta)$ has only a finite number of intersection points.
2. If Q is an intersection point of $\hat{\beta}$, then the inverse image $p^{-1}(Q) \cap \beta$ of Q in β has exactly two points. That is, no more than two distinct points of




Figure 12: *A braid and its projection.*

β are mapped on to any point in $\hat{\beta}$. Such a point Q is said to be a **double point** of $\hat{\beta}$.

3. No vertex of β is mapped to any double point of $\hat{\beta}$.

We say that two regular projections are equivalent if they are the projections of two equivalent braids.

Theorem 2.13.

Every braid has a regular projection.

Proof.

Say we are given a projection $\hat{\beta}$ of our braid β (where $\hat{\beta} = p(\beta)$ as above). Since we have a finite number of polygonal arcs, then to ensure a finite number of intersection points, all we need do is ensure that no two straight lines lie on top of each other. If this does occur, we can correct it by tilting one of the lines as shown in Figure 13. Suppose an intersection point has a pre-image of more than two points. Then we can move subsequent arcs around the intersection point as shown in Figure 14 to ensure that only two arcs meet at the point. Suppose we have a vertex mapped to a double point. Then we merely truncate the vertex as shown in Figure 15, so that the vertex no longer lies on the intersection point. Now, there can only be a finite number of violations of conditions 1-3. And each correctional move we have just defined can be performed so that it does not introduce any more violations. Thus, after a finite number of correctional moves, we have a regular projection. \square

We now alter our regular projections to give us a more intuitive visual perception for our braid. We do this by indicating over/under crossings as in the following definition.




Figure 13: *When two parallel lines in the projection intersect, we can tilt one slightly so that the intersection is now just one point.*




Figure 14: *When more than two lines meet at a point, we can bend all but two of the lines around the point.*




Figure 15: *When a vertex lies on an intersection point, we simply truncate the vertex.*




Figure 16: A double point, and the two possible modifications that can be made to $\hat{\beta}$ to indicate overstrands/understrands.




Figure 17: A regular diagram, re-drawn without any points labeled and with horizontal lines placed at the top and bottom.

Definition 2.14.

Define the projection map onto the x -axis $r : \mathbb{R}^3 \rightarrow \mathbb{R}$ by $r(x, y, z) = x$. Let $\beta \in B_n$, and $\hat{\beta}$ be a regular projection of β . For each double point Q of $\hat{\beta}$ (say an intersection of d_i, d_j), decide which of d_i, d_j is ‘in front’ by seeing which of $r(d_i \cap p^{-1}(Q))$ or $r(d_j \cap p^{-1}(Q))$ is greater. This equates to seeing which of $d_i \cap p^{-1}(Q)$ or $d_j \cap p^{-1}(Q)$ lies closer to the front face of \mathbb{D} . Without loss of generality, we can say d_i is in front. Then remove a small section of $p(d_j)$ around Q in $\hat{\beta}$. This gives us an idea of which strand flows over and which flows under. Figure 16 gives an example of a double point, and the two possible modifications that can be made to $\hat{\beta}$. This modified regular projection is called a **regular diagram** or just **diagram** for the braid β . Two diagrams are said to be **equivalent** if they are the diagrams of equivalent braids.

The idea behind a regular diagram is as follows: We take our braid and lay it out on a page, then draw what we see from above. Crossings are drawn as they appear to us from above, with the over/under strands marked accordingly. The rest is a simple tracing of the braid. From here on, we will draw our regular diagrams without the A_i and B_i labeled, and will place a horizontal line at the top and bottom of the diagram to indicate the top and bottom of the braid. Figure 17 gives an example of this.

3 The braid group

In this section we develop a way to view braids as elements of a group, with a natural operation known as a braid product. We are then able to find a finite presentation of this group.

3.1 Braids as a group

We now move to defining operations on the set of braids as a preliminary step to show that the set of n -braids can be viewed as a group, known as **Artin's n -braid group** or just the **n -braid group**. We begin by introducing a way to multiply braids to give another braid.

Definition 3.1.

Let $\beta_1, \beta_2 \in \mathcal{B}_n$. We shall define a new braid from β_1, β_2 , called the **braid product** of β_1 with β_2 and denoted $\beta_1\beta_2$, as follows: Let β_1, β_2 lie in unit cubes $\mathbb{D}_1, \mathbb{D}_2$ respectively. Identify (glue) the base ($z = 0$) of \mathbb{D}_1 to the top ($z = 1$) of \mathbb{D}_2 . Then scale the union $\mathbb{D}_1 \cup \mathbb{D}_2$ by a factor of $\frac{1}{2}$ and call this \mathbb{D} . Clearly the end points of arcs in β_1 and β_2 are matched up in this identification. The resulting collection of n arcs in \mathbb{D} is denoted $\beta_1\beta_2$, the **braid product** of β_1 with β_2 . Figure 18 gives an example of two 3-braids and their product.

Theorem 3.2.

Let $\beta_1, \beta_2 \in \mathcal{B}_n$. Then $\beta_1\beta_2 \in \mathcal{B}_n$.

Proof.

Let β_1 lie in \mathbb{D}_1 , with braid strings $\{d_1^1, \dots, d_n^1\}$, where each d_i^1 begins at A_i^1 and ends at $B_{j_1(i)}^1$. Let β_2 lie in \mathbb{D}_2 , with braid strings $\{d_1^2, \dots, d_n^2\}$, where each d_i^2 begins at A_i^2 and ends at $B_{j_2(i)}^2$. Then, after we identify the bottom of \mathbb{D}_1 with the top of \mathbb{D}_2 , each B_i^1 is identified with A_i^2 . Thus, for each $1 \leq i \leq n$, the end of d_i^1 (i.e., $B_{j_1(i)}^1$) connects to the start of $d_{j_1(i)}^2$. So $\{d_i^1 \cup d_{j_1(i)}^2 \mid 1 \leq i \leq n\}$ is a set of n polygonal arcs in \mathbb{D} , and we denote each $d_i^1 \cup d_{j_1(i)}^2$ by d_i . In Figure 18 we illustrate this for the braid product of two 3-braids. We now show that the d_i form the braid strings of an n -braid, by direct verification of the definition of an n -braid.

- Since β_1 (respectively β_2) is a braid, then all the d_i^1 (respectively d_i^2) are disjoint. Thus the d_i must be disjoint, each being the union of d_i^1 and $d_{j_1(i)}^2$.
- By definition, each d_i begins at A_i^1 and ends at $B_{j_2(j_1(i))}$.




Figure 18: *Two 3-braids, re-drawn as one above the other, then re-scaled to give their product.*

- Since β_1 (respectively β_2) is a braid, then each level plane in \mathbb{D}_1 (respectively \mathbb{D}_2) intersects each d_i^1 (respectively d_i^2) exactly once. Thus each level plane in \mathbb{D} intersects each d_i exactly once.
- Each d_i^1 (respectively d_i^2) is contained in \mathbb{D}_1 (respectively \mathbb{D}_2). Thus each d_i is contained in \mathbb{D} .

Thus the collection of n arcs d_1, \dots, d_n satisfy the definition of an n -braid, so $\beta_1\beta_2$ is in fact an n -braid. \square

Before we can prove that \mathbf{B}_n forms a group, we require the following lemmas.




Figure 19: Two 3-braids a and b , their product ab , and their product ba .

Lemma 3.3.

Suppose $\beta, \beta', \bar{\beta}, \bar{\beta}' \in \mathcal{B}_n$ with $\beta \sim \beta'$ and $\bar{\beta} \sim \bar{\beta}'$. Then $\beta\bar{\beta} \sim \beta'\bar{\beta}'$

Proof.

We have $\beta \sim \beta'$. Hence there exists the following finite sequence of elementary moves:

$$\beta = \beta_0 \xrightarrow{\Omega_0} \beta_1 \dots \xrightarrow{\Omega_{m-1}} \beta_m = \beta'$$

This induces the following sequence:

$$\beta\bar{\beta} = \beta_0\bar{\beta} \xrightarrow{\Omega_0} \beta_1\bar{\beta} \dots \xrightarrow{\Omega_{m-1}} \beta_m\bar{\beta} = \beta'\bar{\beta}$$

And thus $\beta\bar{\beta} \sim \beta'\bar{\beta}$. We also have $\bar{\beta} \sim \bar{\beta}'$. Hence there exists a finite sequence of elementary moves:

$$\bar{\beta} = \bar{\beta}_0 \xrightarrow{\bar{\Omega}_0} \bar{\beta}_1 \dots \xrightarrow{\bar{\Omega}_{k-1}} \bar{\beta}_k = \bar{\beta}'$$

This induces the following sequence:

$$\beta'\bar{\beta} = \beta'\bar{\beta}_0 \xrightarrow{\bar{\Omega}_0} \beta'\bar{\beta}_1 \dots \xrightarrow{\bar{\Omega}_{k-1}} \beta'\bar{\beta}_k = \beta'\bar{\beta}'$$

And thus $\beta'\bar{\beta} \sim \beta'\bar{\beta}'$. So, since $\beta\bar{\beta} \sim \beta'\bar{\beta}$ and $\beta'\bar{\beta} \sim \beta'\bar{\beta}'$, then (by transitivity of \sim), $\beta\bar{\beta} \sim \beta'\bar{\beta}'$. \square

Lemma 3.4.

Let $\beta_1, \beta_2, \beta_3 \in \mathcal{B}_n$. Then $(\beta_1\beta_2)\beta_3 \sim \beta_1(\beta_2\beta_3)$. That is to say, taking braid products is associative.

Note however that the product of braids is not (in general) commutative. That is, given $\beta, \beta' \in \mathcal{B}_n$, $\beta\beta'$ need not be equivalent to $\beta'\beta$. Figure 19 shows a counter example, giving two 3-braids a and b where ab and ba are not equivalent.

Proof.

Fix diagrams A, B, C for $\beta_1, \beta_2, \beta_3$ respectively, as shown in Figure 20. Then




Figure 20: *Diagrams A, B, C for braids $\beta_1, \beta_2, \beta_3$ respectively.*




Figure 21: *A diagram for $(\beta_1\beta_2)\beta_3$, and a diagram for $\beta_1(\beta_2\beta_3)$.*

Figure 21 a) gives a diagram for $(\beta_1\beta_2)\beta_3$. Similarly, Figure 21 b) gives a diagram for $\beta_1(\beta_2\beta_3)$. But these are identical diagrams. Thus they represent equivalent braids. \square

Definition 3.5.

Let e be the n -braid defined as follows: For each $1 \leq i \leq n$, join A_i to B_i via a straight line segment d_i . The braid e is called the **identity** or **trivial** braid, denoted $\mathbf{1}_n$.

The trivial braid is just that; the simplest n -braid we can find. It is the only n -braid that has a diagram without any crossings. When we braid hair, the trivial braid is what we begin with, before we start adding twists to the strands.

Lemma 3.6.

For any $\beta \in B_n$ we have that $\beta\mathbf{1}_n \sim \beta$ and $\mathbf{1}_n\beta \sim \beta$. Thus $\mathbf{1}_n$ acts as an identity element for braid products.

Proof.

Fix a diagram for β as in Figure 22 a). Now, we know $\mathbf{1}_n$ has a diagram as




Figure 22: Diagrams for β , $\mathbf{1}_n$, and their product $\beta\mathbf{1}_n$.




Figure 23: Contracting the bottom of the vertical strands of a diagram for $\beta\mathbf{1}_n$ to obtain a diagram for β .

shown in Figure 22 b). Thus $\beta\mathbf{1}_n$ has (as one of its diagrams) that of Figure 22 c). However, we can contract the vertical strands at the bottom of the diagram of $\beta\mathbf{1}_n$ as shown in Figure 23, and this is a braid equivalence. Hence $\beta\mathbf{1}_n \sim \beta$ (And an almost identical argument shows that $\mathbf{1}_n\beta \sim \beta$). \square

So attaching the trivial n -braid onto either end of an n -braid β does not change β . This makes sense since all we are effectively doing is lengthening the top or bottom ends of the strands of β .

Lemma 3.7.

Given $\beta \in \mathcal{B}_n$, there exists $\beta' \in \mathcal{B}_n$ such that $\beta\beta' \sim \mathbf{1}_n$ and $\beta'\beta \sim \mathbf{1}_n$. Such a braid β' is called the *inverse* of β , denoted by β^{-1} .

Proof.

Let β be an n -braid in \mathbb{D} . Create a new braid β' by reflecting β in the bottom




Figure 24: *The crossing X and its mirror image over J .*




Figure 25: *Removing the two crossings X and its mirror image.*

face of \mathbb{D} (i.e., the plane $z = 0$, which we shall call J). Then form the n -braid $\beta\beta'$. We now perform elementary moves on $\beta\beta'$ to transform it into the trivial braid as follows: We begin at the junction J of β and β' (i.e., where β ends and β' begins, which is the plane of reflection). We then traverse up the diagram for $\beta\beta'$ (call this B) until we come to the first crossing (i.e., the crossing closest to J). If there happen to be multiple such crossings, choose (any) one. So we have found our first crossing X , say of strands d_i and d_j , and without loss of generality we can assume d_i is the overstrand. Thus we have B looking locally like Figure 24. This is because we know that B has mirror symmetry about J , and there are no other crossings between X and its image below J . Thus some simple elementary moves eliminate the crossing X and its mirror image, as shown in Figure 25. However, we have managed to perform these moves without introducing any new crossings, and still keeping the mirror-symmetry in the diagram. So we now have a new diagram B' for $\beta\beta'$ with two fewer crossings. We then iterate this process on the diagram B' and so forth, each time reducing the number of crossings by 2, and yet preserving the symmetry. Eventually, we are left with an n -braid with no crossings, which is equivalent to the trivial braid. So $\beta\beta' \sim \mathbf{1}_n$. Now, placing β' in \mathbb{D}' and reflecting in the bottom face of \mathbb{D}' gives us β again, hence an almost identical argument shows $\beta'\beta \sim \mathbf{1}_n$. \square

So as we can see, it is extremely easy to construct the inverse β^{-1} of an n -braid β . The existence of such a braid means that we can undo any braid from below (or above). This is in contrast to knot theory, where it is known that the connect sum of any two non-trivial knots will give another non-trivial knot. See [8] for more details about the connect-sum of knots.

Our final step is to define a set in which two equivalent braids are considered the same. This quotient of \mathcal{B}_n will end up being our group of n -braids.

Definition 3.8.

Let $\beta \in \mathcal{B}_n$. We denote the \sim -equivalence class of β by $[\beta]$. Denote by \mathbf{B}_n the set of all \sim -equivalence classes of n -braids. That is,

$$\mathbf{B}_n := \mathcal{B}_n / \sim$$

Definition 3.9.

Let $[\beta], [\beta'] \in \mathbf{B}_n$. Define an operation on equivalence classes of braids as follows:

$$[\beta] \cdot [\beta'] := [\beta\beta']$$

where $\beta\beta'$ is the braid product of β and β' .

Theorem 3.10.

*The set \mathbf{B}_n forms a group, with operation \cdot as defined above. This group is called the **n -braid group** or **Artin's n -braid group**.*

Proof.

Let $[\beta_1], [\beta_2], [\beta_3] \in \mathbf{B}_n$.

1.

$[\beta_1] \cdot [\beta_2] = [\beta_1\beta_2] \in \mathbf{B}_n$ by Theorem 3.2.

So \mathbf{B}_n is closed under the operation \cdot .

2.

$$\begin{aligned} ([\beta_1] \cdot [\beta_2]) \cdot [\beta_3] &= [\beta_1\beta_2] \cdot [\beta_3] \\ &= [(\beta_1\beta_2)\beta_3] \\ &= [\beta_1(\beta_2\beta_3)] \text{ by Lemma 3.4} \\ &= [\beta_1] \cdot [\beta_2\beta_3] \\ &= [\beta_1] \cdot ([\beta_2] \cdot [\beta_3]) \end{aligned}$$

So \cdot is an associative operation on \mathbf{B}_n .



Figure 26: A diagram for σ_i , and for σ_i^{-1} .

3.

$[\mathbf{1}_n] \cdot [\beta_1] = [\mathbf{1}_n \beta_1] = [\beta_1]$ by Lemma 3.6.

$[\beta_1] \cdot [\mathbf{1}_n] = [\beta_1 \mathbf{1}_n] = [\beta_1]$ by Lemma 3.6.

So $[\mathbf{1}_n]$ is the identity element for \mathbf{B}_n .

4.

$[\beta_1] \cdot [\beta_1^{-1}] = [\beta_1 \beta_1^{-1}] = [\mathbf{1}_n]$ by Lemma 3.7

$[\beta_1^{-1}] \cdot [\beta_1] = [\beta_1^{-1} \beta_1] = [\mathbf{1}_n]$ by Lemma 3.7

So $[\beta_1^{-1}]$ is the inverse to $[\beta_1]$, written $[\beta_1]^{-1}$

Thus (\mathbf{B}_n, \cdot) satisfies the definition of a group. □

Note: To avoid awkward notation, we will denote an n -braid $\beta \in \mathcal{B}_n$ and its equivalence class $[\beta] \in \mathbf{B}_n$ both by β (since we intuitively hold all equivalent braids as being the same). It should be obvious in the context of what is being said which definition is being assumed.

3.2 A presentation of the braid group

We now move on to finding an explicit presentation for \mathbf{B}_n . As it turns out, \mathbf{B}_n is finitely presented, and we shall find one such finite presentation.

Definition 3.11.

For $1 \leq i \leq n - 1$, define the n -braid σ_i as the braid represented by the diagram in Figure 26 a). That is, σ_i is the braid with only one crossing, where the string from A_i to B_{i+1} crosses under the string from A_{i+1} to B_i . The inverse n -braid of σ_i (denoted σ_i^{-1} as usual) is thus given as the braid represented by the diagram in Figure 26 b). That is, σ_i is the braid with only one crossing, where the string from A_i to B_{i+1} crosses over the string from A_{i+1} to B_i . The set $\{\sigma_1, \dots, \sigma_{n-1}\}$ is known as the set of **Artin generators** for the braid group \mathbf{B}_n .

The σ_i 's are the simplest n -braids (after the trivial braid), having diagrams with only one crossing (necessarily between two adjacent strands). It should not be




Figure 27: A regular diagram, re-drawn so that the two crossings that were originally at the same height are no longer.

hard to see that any n -braid can be partitioned into the σ'_i s and their inverses, and we will see shortly that these actually form a generating set for \mathbf{B}_n .

Theorem 3.12.

Let $\beta \in \mathbf{B}_n$. Then β can be written as the product of the σ'_i s and their inverses (i.e., $\beta = \sigma_{i_1}^{\epsilon_1} \dots \sigma_{i_k}^{\epsilon_k}$ with $1 \leq i_1, \dots, i_k \leq n - 1$, $\epsilon_i \in \{\pm 1\}$ for $1 \leq i \leq k$, $k \in \mathbb{N}$).

Proof.

Given an n -braid β , we know that there exists a diagram for β . This diagram has a finite number of crossings. By shifting the crossings up or down slightly, we can have a diagram for β whereby no two crossings are at the same height, and we give an example of this in Figure 27. We can then separate the crossings by level planes, and thus partition β into sections, with each section having only one crossing in its diagram. We can then perform a further set of elementary moves to transform each section into a braid in which all but two (adjacent) strands are entirely vertical. Thus we have partitioned our braid as the product of the σ_i and their inverses. Figure 28 gives an example of this partitioning and subsequent straightening. Hence we can use a finite number of level planes to partition β into sections, each section being equivalent to σ_i^ϵ for some $1 \leq i \leq n - 1$ and $\epsilon \in \{\pm 1\}$. Thus β is equivalent to the finite product of these (i.e., $\beta = \sigma_{i_1}^{\epsilon_1} \dots \sigma_{i_k}^{\epsilon_k}$). \square

We now describe some relations between the Artin generators, which are immediately obvious as they are consequences of some simple elementary moves. What is not so obvious however is that the relations we are about to state form a set of defining relations in a presentation for \mathbf{B}_n .




Figure 28: A braid with level planes between each pair of crossings, re-drawn so that most of the strands are vertical lines.




Figure 29: A diagram for $\sigma_i\sigma_j$, and after performing the dotted elementary move we have a diagram for $\sigma_j\sigma_i$.

Theorem 3.13.

The following relations hold in \mathbf{B}_n :

1. $\sigma_i\sigma_j = \sigma_j\sigma_i$ for any $1 \leq i, j \leq n - 1$ with $|i - j| \geq 2$
2. $\sigma_i\sigma_{i+1}\sigma_i = \sigma_{i+1}\sigma_i\sigma_{i+1}$ for any $1 \leq i \leq n - 2$

Proof.

1. In Figure 29 a) we have a diagram for $\sigma_i\sigma_j$ for some $1 \leq i, j \leq n - 1$ with $|i - j| \geq 2$. We can also see a dotted elementary move drawn in on this diagram. Figure 29 b) is the resultant diagram after this elementary move is performed, and is a diagram for $\sigma_j\sigma_i$.

2. In Figure 30 a) we have a diagram for $\sigma_i\sigma_{i+1}\sigma_i$ for some $1 \leq i \leq n - 2$, with




Figure 30: A diagram for $\sigma_i\sigma_{i+1}\sigma_i$, and after performing the dotted elementary move we have a diagram for $\sigma_{i+1}\sigma_i\sigma_{i+1}$.

a dotted elementary move drawn in (technically, this is the composition of some number of elementary moves, which should be obvious to see). Figure 30 b) is the resultant diagram after this elementary move is performed, and is a diagram for $\sigma_{i+1}\sigma_i\sigma_{i+1}$. \square

Theorem 3.14.

For any $n \geq 1$, the n -braid group \mathbf{B}_n has a presentation given by

$$\mathbf{B}_n = \langle \sigma_1, \dots, \sigma_{n-1} \mid \begin{aligned} &\sigma_i\sigma_j = \sigma_j\sigma_i \text{ for } 1 \leq i, j \leq n-1 \text{ and } |i-j| \geq 2, \\ &\sigma_i\sigma_{i+1}\sigma_i = \sigma_{i+1}\sigma_i\sigma_{i+1} \text{ for } 1 \leq i \leq n-2 \end{aligned} \rangle$$

Proof.

We begin by defining a group G by

$$G := \langle x_1, \dots, x_{n-1} \mid \begin{aligned} &x_i x_j = x_j x_i \text{ for } 1 \leq i, j \leq n-1 \text{ and } |i-j| \geq 2, \\ &x_i x_{i+1} x_i = x_{i+1} x_i x_{i+1} \text{ for } 1 \leq i \leq n-2 \end{aligned} \rangle$$

Theorem 3.12 shows that the Artin generators $\{\sigma_1, \dots, \sigma_{n-1}\}$ form a generating set for \mathbf{B}_n . So we define a map $\bar{\phi} : \{x_1, \dots, x_{n-1}\} \rightarrow \mathbf{B}_n$ via $\bar{\phi}(x_i) = \sigma_i$. This extends to a map $\phi : G \rightarrow \mathbf{B}_n$. Furthermore, we have that

1. If $|i-j| > 1$, then

$$\begin{aligned} \phi(x_i x_j) &= \sigma_i \sigma_j \\ &= \sigma_j \sigma_i \text{ since } |i-j| > 1 \text{ (by Theorem 3.13)} \\ &= \phi(x_j x_i) \end{aligned}$$




Figure 31: The diagrams C_{σ_i} and $C_{\sigma_i^{-1}}$ of σ_i and σ_i^{-1} respectively.




Figure 32: Concatenating the diagrams C_1^{-1} , C_3 and C_2 to give the diagram $C_1^{-1}C_3C_2$.

2. For any $1 \leq i < n - 2$ we have

$$\begin{aligned}
 \phi(x_i x_{i+1} x_i) &= \sigma_i \sigma_{i+1} \sigma_i \\
 &= \sigma_{i+1} \sigma_i \sigma_{i+1} \text{ since } 1 \leq i \leq n - 2 \text{ (By Theorem 3.13)} \\
 &= \phi(x_{i+1} x_i x_{i+1})
 \end{aligned}$$

Thus ϕ is a homomorphism. And clearly each σ_i lies in the image of ϕ , since $\phi(x_i) = \sigma_i$ for any $1 \leq i \leq n - 1$. So ϕ is surjective.

To prove the injectivity of ϕ , we proceed as follows: We wish to show that $\text{Ker}(\phi) = \{\mathbf{1}\}$. So fix an element $g \in \text{Ker}(\phi)$, and fix a word $w = x_{i_1}^{\epsilon_1} \dots x_{i_k}^{\epsilon_k}$ representing g . Thus $\phi(g) = \phi([w]) = \sigma_{i_1}^{\epsilon_1} \dots \sigma_{i_k}^{\epsilon_k} = \mathbf{1}_{\mathbf{B}_n}$ (the identity in \mathbf{B}_n) since $g \in \text{Ker}(\phi)$. We form a representative braid $\beta \in \mathcal{B}_n$ for $\phi(g)$ as follows: Define a diagram C_{σ_i} for each σ_i in the obvious manner, and similarly define a diagram $C_{\sigma_i^{-1}}$ for σ_i^{-1} . See Figure 31 for examples of these diagrams. Now, form a diagram D^w by concatenating $C_{\sigma_{i_1}^{\epsilon_1}} \dots C_{\sigma_{i_k}^{\epsilon_k}}$ (i.e., by drawing $C_{\sigma_{i_1}^{\epsilon_1}}$, and then $C_{\sigma_{i_2}^{\epsilon_2}}$ directly below it, and so on). Note that this definition extends to *any* word w over the Artin generators representing an element in G . Figure 32 gives an example of the diagram $D^{\sigma_1^{-1} \sigma_3 \sigma_2}$, formed by concatenating C_1^{-1} , C_3 and C_2 to give the diagram $C_1^{-1}C_3C_2$.

We denote by $\beta^w \in \mathcal{B}_n$ the n -braid represented by the diagram D^w (again extending this definition to any word w representing an element in G). Clearly, β^w

is a representative for $\phi(g)$ since $[\beta^w] = \sigma_{i_1}^{\epsilon_1} \dots \sigma_{i_k}^{\epsilon_k}$. It is important here to realise that if two diagrams D^w and $D^{w'}$ are identical (not just equivalent), then w and w' must be identical as words. Now, since $\phi(g) = \mathbf{1}_{\mathbf{B}_n}$, we have that $[\beta^w] = \mathbf{1}_{\mathbf{B}_n}$ (i.e., $\beta^w \sim \mathbf{1}_n$). So there exists a finite sequence of l elementary moves and their inverses $\Omega_1^{m_1}, \dots, \Omega_l^{m_l}$ (where $m_j \in \{\pm 1\}$ for all $1 \leq j \leq l$) that transforms β^w into $\mathbf{1}_n$. So we end up with

$$\beta^w = \beta_0 \xrightarrow{\Omega_1^{m_1}} \beta_1 \xrightarrow{\Omega_2^{m_2}} \dots \xrightarrow{\Omega_l^{m_l}} \beta_l = \mathbf{1}_n$$

where all the β_i are intermediate n -braids. We can view the elementary moves as acting on the diagrams for each intermediate braid as described above. So, defining D_i as the diagram for β_i for each i , and denoting an elementary move on a diagram by $\bar{\Omega}^{\pm 1}$, we obtain the following sequence:

$$D^w = D_0 \xrightarrow{\bar{\Omega}_1^{m_1}} D_1 \xrightarrow{\bar{\Omega}_2^{m_2}} \dots \xrightarrow{\bar{\Omega}_l^{m_l}} D_l = D^1$$

where D^1 is the trivial diagram for the trivial braid $\mathbf{1}_n$ (i.e., n vertical lines). Figure 33 gives an example of this diagram. Our claim is that, given two words w_1 and w_2 over the x_i , if D^{w_1} and D^{w_2} differ by some elementary move $\bar{\Omega}$ (i.e., $D^{w_1} \xrightarrow{\bar{\Omega}} D^{w_2}$), then $[w_1] =_G [w_2]$. We show this by dealing with the finite number of possible cases individually.

So, given $w_1 = x_{\tau_1}^{\alpha_1} \dots x_{\tau_y}^{\alpha_y}$ and $w_2 = x_{\rho_1}^{\gamma_1} \dots x_{\rho_z}^{\gamma_z}$, we begin with the diagram D^{w_1} , and perform the elementary move $\bar{\Omega}^s$ ($s \in \{\pm 1\}$) which acts on some braid string d and transforms D^{w_1} into D^{w_2} . Thus we have a triangle ΔABC in the plane, and we either replace $AC \cup BC$ with AB (if $s = -1$), or replace AB with $AC \cup BC$ (if $s = +1$). We must look at what could possibly lie inside the triangle ΔABC (where here the triangle is actually projected in the plane; obviously the triangle itself must be empty in 3-space by definition of an elementary move).

Case 1. ΔABC is empty, as in Figure 34.

Then the word w_1 also has as its diagram D^{w_2} , since we have not moved, created nor destroyed any crossings. Thus $w_1 = w_2$, and so $[w_1] =_G [w_2]$.

Case 2. ΔABC contains part of one other strand d' , where d' enters and exits ΔABC under AB , as in Figure 35.

If $s = 1$, then we have (for some i) removed the sub-diagram $C_i C_i^{-1}$ from somewhere in our diagram, and thus we now have the diagram of $x_{\tau_1}^{\alpha_1} \dots x_{\tau_{j-1}}^{\alpha_{j-1}} x_{\tau_{j+2}}^{\alpha_{j+2}} \dots x_{\tau_y}^{\alpha_y}$ (where $x_{\tau_j}^{\alpha_j} = x_i$ and $x_{\tau_{j+1}}^{\alpha_{j+1}} = x_i^{-1}$), which must be identical to the word w_2 . If $s = -1$, then we have (for some i) added the sub-diagram $C_i C_i^{-1}$ somewhere in our diagram, and thus we now have the diagram of $x_{\tau_1}^{\alpha_1} \dots x_{\tau_j}^{\alpha_j} x_i x_i^{-1} x_{\tau_{j+1}}^{\alpha_{j+1}} \dots x_{\tau_y}^{\alpha_y}$, which must be identical to the word w_2 . But for either instance, all we have done is perform a free cancellation (or its inverse) on w_1 , and thus $[w_1] =_G [w_2]$.




Figure 33: *The diagram D^1 for the trivial braid $\mathbf{1}_n$.*




Figure 34: *An elementary move on d , where $\triangle ABC$ is empty.*




Figure 35: *An elementary move on d , where $\triangle ABC$ contains part of one other strand d' that enters and exits $\triangle ABC$ under AB .*

Case 3. ΔABC contains part of one other strand d' , where d' enters ΔABC under AC , and exits under BC , as in Figure 36.

If $s = 1$, then we have (for some i) added the sub-diagram $C_i^{-1}C_i$ somewhere in our diagram, and thus we now have the diagram of $x_{\tau_1}^{\alpha_1} \dots x_{\tau_j}^{\alpha_j} x_i^{-1} x_i x_{\tau_{j+1}}^{\alpha_{j+1}} \dots x_{\tau_y}^{\alpha_y}$, which must be identical to the word w_2 . If $s = -1$, then we have (for some i) removed the sub-diagram $C_i^{-1}C_i$ from somewhere in our diagram, and thus we now have the diagram of $x_{\tau_1}^{\alpha_1} \dots x_{\tau_{j-1}}^{\alpha_{j-1}} x_{\tau_{j+2}}^{\alpha_{j+2}} \dots x_{\tau_y}^{\alpha_y}$ (where $x_{\tau_j}^{\alpha_j} = x_i^{-1}$ and $x_{\tau_{j+1}}^{\alpha_{j+1}} = x_i$), which must be identical to the word w_2 . But for either instance, all we have done is perform a free cancellation (or its inverse) on w_1 , and thus $[w_1] =_G [w_2]$.

Case 4. ΔABC contains part of one other strand d' , where d' enters ΔABC under one of AB or $AC \cup BC$ and exits under the other, as in Figure 37.

Regardless of the sign of s , we have (for some i and some $\epsilon \in \{\pm 1\}$), shifted the sub-diagram C_i^ϵ up or down the order in our diagram, past some number of the other $C_r^{\pm 1}$ with $r \neq i - 1, i, i + 1$. Thus we now have either the diagram of $x_{\tau_1}^{\alpha_1} \dots x_{\tau_j}^{\alpha_j} x_i^\epsilon x_{\tau_{j+1}}^{\alpha_{j+1}} \dots x_{\tau_{v-1}}^{\alpha_{v-1}} x_{\tau_{v+1}}^{\alpha_{v+1}} \dots x_{\tau_y}^{\alpha_y}$ if C_i^ϵ is shifted up (where $x_{\tau_v}^{\alpha_v} = x_i^\epsilon$ and $|\tau_q - i| \geq 2$ for all $j + 1 \leq q \leq v - 1$). Or the diagram of $x_{\tau_1}^{\alpha_1} \dots x_{\tau_{v-1}}^{\alpha_{v-1}} x_{\tau_{v+1}}^{\alpha_{v+1}} \dots x_{\tau_j}^{\alpha_j} x_i^\epsilon x_{\tau_{j+1}}^{\alpha_{j+1}} \dots x_{\tau_y}^{\alpha_y}$ if C_i^ϵ is shifted down (where $x_{\tau_v}^{\alpha_v} = x_i^\epsilon$ and $|\tau_q - i| \geq 2$ for all $v + 1 \leq q \leq j$). Either way, the resultant word must be identical to the word w_2 . But all we have done is swapped $x_i^\epsilon x_j^{\pm 1}$ with $x_j^{\pm 1} x_i^\epsilon$ (or vice versa) for $|i - j| \geq 2$ some number of times in w_1 , which is an equivalence in G since $\sigma_i \sigma_j = \sigma_j \sigma_i$. Thus $[w_1] =_G [w_2]$.

Case 5. ΔABC contains part of two strands d', d'' , with one crossing between them, where both strands of the crossing enter ΔABC under AC and exit under BC , as in Figure 38.

If $s = 1$ then, for some i , we have replaced the sub-diagram C_{i+1}^ϵ with the sub-diagram $C_i^{-1} C_{i+1}^{-1} C_i^\epsilon C_{i+1} C_i$. Thus we now have the diagram of $x_{\tau_1}^{\alpha_1} \dots x_{\tau_j}^{\alpha_j} x_i^{-1} x_{i+1}^{-1} x_i^\epsilon x_{i+1} x_i x_{\tau_{j+2}}^{\alpha_{j+2}} \dots x_{\tau_y}^{\alpha_y}$ (where $x_{\tau_{j+1}}^{\alpha_{j+1}} = x_{i+1}^\epsilon$), which must be identical to the word w_2 . If $s = -1$, then, for some i , we have replaced the sub-diagram $C_i^{-1} C_{i+1}^{-1} C_i^\epsilon C_{i+1} C_i$ with the sub-diagram C_{i+1}^ϵ . Thus we now have the diagram of $x_{\tau_1}^{\alpha_1} \dots x_{\tau_j}^{\alpha_j} x_{i+1}^\epsilon x_{\tau_{j+6}}^{\alpha_{j+6}} \dots x_{\tau_y}^{\alpha_y}$, (where $x_i^{-1} x_{i+1}^{-1} x_i^\epsilon x_{i+1} x_i$ and $x_{\tau_{j+1}}^{\alpha_{j+1}} \dots x_{\tau_{j+5}}^{\alpha_{j+5}}$ are identical as words), which must be identical to the word w_2 . But for either instance, all we have done is interchange x_{i+1}^ϵ with $x_i^{-1} x_{i+1}^{-1} x_i^\epsilon x_{i+1} x_i$ (or vice versa), which is an equivalence in G since $\sigma_i \sigma_{i+1} \sigma_i = \sigma_{i+1} \sigma_i \sigma_{i+1}$. Thus $[w_1] =_G [w_2]$.

Case 6. ΔABC contains one crossing, where the strands of the crossing enter and exit ΔABC under AC, BC, AB in an arbitrary manner.

It is obvious that we can break up $\bar{\Omega}^s$ into an initial series of cases 1-4, after which the strands of the crossing enter ΔABC under AC and exit under BC . We then apply case 5, so we still conclude that $[w_1] =_G [w_2]$.

Note that for the above cases all strands enter and exit ΔABC as understrands,




Figure 36: An elementary move on d , where $\triangle ABC$ contains part of one other strand d' that enters $\triangle ABC$ under AC and exits under BC .




Figure 37: An elementary move on d , where $\triangle ABC$ contains part of one other strand d' , where d' enters $\triangle ABC$ under one of $AB, AC \cup BC$ and exits under the other.




Figure 38: An elementary move on d , where $\triangle ABC$ contains part of two strands with one crossing between them, where both strands of the crossing enter $\triangle ABC$ under AC and exit under BC .

and that the point C is always to the right of AB . A set of very similar arguments show that $[w_1] =_G [w_2]$ when strands enter and exit ΔABC as overstrands, and/or when C is to the left of AB .

Case 7. ΔABC contains parts of arbitrarily many strands, that have arbitrarily many intersections between them.

Then we can break up $\bar{\Omega}^s$ into a finite series of the above cases, where at each intermediate stage the intermediate words are still equivalent in G . So we still have $[w_1] =_G [w_2]$.

Thus, if $\bar{\Omega}^s$ is an elementary move transforming D^{w_1} into D^{w_2} , then $[w_1] =_G [w_2]$. So, to complete the proof that ϕ is injective, recall from the beginning of our proof that we had a finite sequence of elementary moves transforming D^w into D^1 . Thus we must have that $[w] =_G [1]$ (i.e., $g =_G \mathbf{1}_G$). But g was an arbitrary element of $\text{Ker}(\phi)$. Thus $\text{Ker}(\phi) = \{\mathbf{1}_G\}$, and so ϕ is injective. Thus ϕ is an isomorphism, so \mathbf{B}_n has a presentation identical to that of G . \square

Given that we now have a (finite) presentation of \mathbf{B}_n , we will give a lot of results that can be derived purely from the presentation of \mathbf{B}_n , without any knowledge of the geometry behind our construction. However, where possible we will attempt to support these results with pictures, to ensure that we maintain an intuitive feel for the subject without getting lost in the algebra.

4 Properties of the braid group

Now that we have established a group \mathbf{B}_n associated with braids, and have a finite presentation for this group, we are able to look at properties of this group and thus extract more information about braids and their structure. We can use this group to find braids that commute with all other braids, or see what happens when we add more strings to an existing braid, or develop methods of distinguishing braids. We can even find interesting subgroups of \mathbf{B}_n that have a nice geometric interpretation.

4.1 Some results about the braid group

In this section we discuss some simple properties of the braid group \mathbf{B}_n , that come primarily from the presentation we have just found. We will discuss the centre of \mathbf{B}_n , and see how to view \mathbf{B}_m as a subgroup of \mathbf{B}_n for $m < n$. This will allow us to calculate the size of \mathbf{B}_n .

Theorem 4.1.

For $n = 2$, $Z(\mathbf{B}_n) = \mathbf{B}_n = \langle \sigma_1 | - \rangle$. For $n \geq 3$, the centre of \mathbf{B}_n is generated by the element Δ_n^2 , where $\Delta_n = (\sigma_1 \dots \sigma_{n-1})(\sigma_1 \dots \sigma_{n-2}) \dots (\sigma_1 \sigma_2)(\sigma_1)$. That is, $Z(\mathbf{B}_n) = \langle \Delta_n^2 \rangle$.

Proof.

The case $n = 2$ is trivial, since $\mathbf{B}_2 = \langle \sigma_1 | - \rangle$. For $n \geq 2$, we will later prove in Theorem 5.10 that $\Delta_n^2 \in Z(\mathbf{B}_n)$. For a full proof that Δ_n^2 generates $Z(\mathbf{B}_n)$, see [7], pp 246-247. \square

Figure 39 shows diagrams for Δ_5 and Δ_5^2 , and helps illustrate the simple geometric interpretation of Δ_n in general. It is constructed from the trivial braid $\mathbf{1}_n$ as follows: Grab the bottom end of $\mathbf{1}_n$ and give it a half-twist anti-clockwise (when looking up the braid from bottom to top). This twist permutes the bottom end-points of the braid by sending B_1 to B_{n-1} , B_2 to B_{n-2} , and so on. It is not hard to see that $\Delta_n \sigma_i = \sigma_{n-i} \Delta_n$ for any $1 \leq i \leq n-1$; just push σ_i up through Δ_n , and when it comes out at the top it has been given a half twist and shifted across. We will later show this in section 5.1. Similarly, Δ_n^2 is merely $\mathbf{1}_n$ with a full-twist put in. So if we put a full twist in at the bottom of any braid β , we can then give the braid itself a full twist so that the full twist at the bottom disappears, and instead we have a full twist at the top. This is the geometric essence behind the fact that Δ_n^2 lies in $Z(\mathbf{B}_n)$.




Figure 39: *The braid Δ_5 , and its square Δ_5^2 .*




Figure 40: A diagram B of an m -braid, and a diagram of its image under ψ . Note the $n - m$ vertical lines added to the right of B to give the image under ψ .

We now move on to discuss the relationship between different braid groups. What may not seem obvious is that, for any $1 \leq m \leq n$, \mathbf{B}_m is isomorphic to a subgroup of \mathbf{B}_n . We will give an algebraic proof here, but follow it up with a geometric interpretation which is much more intuitive.

Theorem 4.2.

Let $1 \leq m \leq n$. Define a map ψ by

$$\begin{aligned} \psi : \mathbf{B}_m &\rightarrow \mathbf{B}_n \\ \psi(\sigma_{i_1}^{\epsilon_1} \dots \sigma_{i_k}^{\epsilon_k}) &= \acute{\sigma}_{i_1}^{\epsilon_1} \dots \acute{\sigma}_{i_k}^{\epsilon_k} \end{aligned}$$

where $\{\sigma_1, \dots, \sigma_{m-1}\}$ are the generators of \mathbf{B}_m and $\{\acute{\sigma}_1, \dots, \acute{\sigma}_{n-1}\}$ are the generators of \mathbf{B}_n . Then ψ defines a homomorphism from \mathbf{B}_m to \mathbf{B}_n . Also, ψ is injective, and thus \mathbf{B}_m can be viewed as a subgroup of \mathbf{B}_n .

Another way to look at the image $\psi(\beta)$ of β is to imagine $n - m$ straight vertical strings being added to the right hand side of a diagram for β (B), as shown in Figure 40. Thus the inverse of ψ (acting on the set $\psi(\mathbf{B}_m)$) can immediately be seen as removing the $n - m$ straight vertical strings from the right hand side of $\psi(\beta)$ to obtain β once again.

Proof.

To check ψ is a homomorphism, we need only check that the relators in a presentation of \mathbf{B}_m hold under ψ . So let $1 \leq i \leq m - 2$.

$$\begin{aligned} \Rightarrow \psi(\sigma_i \sigma_{i+1} \sigma_i) &= \acute{\sigma}_i \acute{\sigma}_{i+1} \acute{\sigma}_i \\ &= \acute{\sigma}_{i+1} \acute{\sigma}_i \acute{\sigma}_{i+1} \\ &= \psi(\sigma_{i+1} \sigma_i \sigma_{i+1}) \end{aligned}$$

Now let $1 \leq i, j \leq n - 1$ with $|i - j| \geq 2$.

$$\begin{aligned} \Rightarrow \psi(\sigma_i \sigma_j) &= \acute{\sigma}_i \acute{\sigma}_j \\ &= \acute{\sigma}_j \acute{\sigma}_i \\ &= \psi(\sigma_j \sigma_i) \end{aligned}$$

So the image under ψ of every relation in \mathbf{B}_m holds in \mathbf{B}_n . Now suppose $\psi(\beta) = \mathbf{1}_n$. Then, viewing $\psi(\beta)$ and $\mathbf{1}_n$ as elements of \mathcal{B}_n , there is a finite sequence of elementary moves

$$\psi(\beta) = \beta_0 \xrightarrow{\Omega_0} \beta_1 \xrightarrow{\Omega_1} \dots \xrightarrow{\Omega_{k-1}} \beta_k = \mathbf{1}_n$$

where β_i is an n -braid, and β_{i+1} is obtained from β_i by applying the elementary move Ω_i . Now, we remove the final $n - m$ strings from each β_i to form $\bar{\beta}$ (note that these strings need not be straight any more, since we have applied some elementary moves). So this induces a sequence of elementary moves

$$\beta = \bar{\beta}_0 \xrightarrow{\bar{\Omega}_0} \bar{\beta}_1 \xrightarrow{\bar{\Omega}_1} \dots \xrightarrow{\bar{\Omega}_{k-1}} \bar{\beta}_k = \mathbf{1}_m$$

where $\bar{\Omega}_i$ is the identity move if Ω_i acts on the l^{th} braid string of β_i with $m \leq l \leq n$, and is the elementary move Ω_i (acting just on the braid strings of $\bar{\beta}_i$) otherwise. Thus $\beta \sim \mathbf{1}_m$, and so ψ is injective. \square

Corollary 4.3.

Let $\beta \in \mathcal{B}_m$ be such that $\beta \approx \mathbf{1}_m$. Let $[\beta'] = \psi([\beta]) \in \mathbf{B}_n$, with $1 \leq m \leq n$ (ψ as defined previously). Then $\beta' \approx \mathbf{1}_n$.

Proof.

This is immediate from the injectivity of ψ . \square

That is to say, if we add any number of straight arcs to the end of a non-trivial braid, we end up with another non-trivial braid.

We can now investigate what happens to a braid if we put an arbitrary number of twists between any two adjacent strings, and thus show something about the size of \mathbf{B}_n .

Theorem 4.4.

Let σ_i be an Artin generator for \mathbf{B}_n . Then $|\sigma_i| = \infty$. That is, for any $k > 0$, we have $\sigma_i^k \neq \mathbf{1}_n$.

Proof.

$\mathbf{B}_2 = \langle \acute{\sigma}_1 | - \rangle$ is a free group (where the dash is to help with notation later on). Thus $|\acute{\sigma}_1| = \infty$. Now, since $\psi : \mathbf{B}_2 \rightarrow \mathbf{B}_n$ is an injective homomorphism for any $n \geq 2$, then $|\psi(\acute{\sigma}_1)| = \infty$. But $\psi(\acute{\sigma}_1) = \sigma_1$, the first Artin generator for \mathbf{B}_n . So $|\sigma_1| = \infty$. We now proceed by induction to derive a contradiction, so suppose $|\sigma_i| = l$ for some integer $l \geq 1$ and some $2 \leq i \leq n - 1$ (where σ_i is an Artin generator for \mathbf{B}_n). Then we have $(\sigma_{i-1}\sigma_i\sigma_{i-1})\sigma_i^l(\sigma_{i-1}\sigma_i\sigma_{i-1})^{-1} = \mathbf{1}_n$. But

$$\begin{aligned} \sigma_{i-1}\sigma_i^l\sigma_{i-1}^{-1} &= (\sigma_{i-1}\sigma_i\sigma_{i-1}^{-1})^l \\ &= (\sigma_i^{-1}\sigma_{i-1}\sigma_i)^l \\ &= \sigma_i^{-1}\sigma_{i-1}^l\sigma_i \end{aligned}$$

Thus we have

$$\begin{aligned} \mathbf{1}_n &= (\sigma_{i-1}\sigma_i\sigma_{i-1})\sigma_i^l(\sigma_{i-1}\sigma_i\sigma_{i-1})^{-1} \\ &= \sigma_{i-1}\sigma_i(\sigma_{i-1}\sigma_i^l\sigma_{i-1}^{-1})\sigma_i^{-1}\sigma_{i-1}^{-1} \\ &= \sigma_{i-1}\sigma_i(\sigma_i^{-1}\sigma_{i-1}^l\sigma_i)\sigma_i^{-1}\sigma_{i-1}^{-1} \\ &= \sigma_{i-1}\sigma_{i-1}^l\sigma_{i-1}^{-1} \\ &= \sigma_{i-1}^l \end{aligned}$$

So if $\sigma_i^l = \mathbf{1}_n$ for some integer $l \geq 1$ and some $2 \leq i \leq n - 1$, then $\sigma_j^l = \mathbf{1}_n$ for all $j \leq i$. More specifically, $\sigma_1^l = \mathbf{1}_n$, which is a contradiction since $|\sigma_1| = \infty$. \square

Corollary 4.5.

$|\mathbf{B}_n| = \infty$ for any $n \geq 2$.

Proof.

If $n \geq 2$, then \mathbf{B}_n has at least one generator σ_1 , and this generator has infinite order. \square

It is not surprising that \mathbf{B}_n is an infinite group for any $n \geq 2$, since applying an arbitrary number of twists between adjacent strings will give us a braid of arbitrary complexity. So it turns out we indeed have many braids to deal with. However, as we shall see in section 5.1, there is an algorithm that checks if two arbitrary braids are equivalent or not. This ensures we have a way of completely classifying all braids.

4.2 Braid invariants

We now move on to defining some homomorphisms between \mathbf{B}_n and other well-known groups. The first of these is a homomorphism onto a finite group, with a simple geometric interpretation.

Definition 4.6.

Let $\beta \in \mathbf{B}_n$. Suppose the i^{th} braid string d_i joins A_i to $B_{j(i)}$ for $1 \leq i \leq n$. Define the **braid permutation** $\pi : \mathbf{B}_n \rightarrow \mathbf{S}_n$, where \mathbf{S}_n is the symmetric group on n elements, by

$$\pi(\beta) := \begin{pmatrix} 1 & 2 & \dots & n \\ j(1) & j(2) & \dots & j(n) \end{pmatrix}$$

So for each n -braid β , we can associate a unique permutation of the numbers $1, 2, \dots, n$ with β , that corresponds to the way the strings of β permute the starting points with the end points. However, the image under π of an n -braid β only encodes part of the information of the braid. It is not hard to see that, given some $\rho \in \mathbf{S}_n$, there are in fact many different n -braids that map to ρ via π .

Theorem 4.7.

The braid permutation π as defined above is a homomorphism.

Proof.

Clearly $\pi(\sigma_i) = (i \ i+1)$ for any $1 \leq i \leq n-1$. We proceed to check that $\pi(r) = (1)$ for any relator in the standard presentation of \mathbf{B}_n , where (1) is the identity element in \mathbf{S}_n . So let $1 \leq i \leq n-2$. Then

$$\begin{aligned} \pi(\sigma_i \sigma_{i+1} \sigma_i) &= \pi(\sigma_i) \pi(\sigma_{i+1}) \pi(\sigma_i) \\ &= (i \ i+1)(i+1 \ i+2)(i \ i+1) \\ &= (i \ i+2) \end{aligned}$$

We also have

$$\begin{aligned} \pi(\sigma_{i+1} \sigma_i \sigma_{i+1}) &= (i+1 \ i+2)(i \ i+1)(i+1 \ i+2) \\ &= (i \ i+2) \end{aligned}$$

Thus $\pi(\sigma_i \sigma_{i+1} \sigma_i) = \pi(\sigma_{i+1} \sigma_i \sigma_{i+1})$. Now let $1 \leq i \leq n-1$ with $|i-j| \geq 2$. Then

$$\begin{aligned} \pi(\sigma_i \sigma_j) &= \pi(\sigma_i) \pi(\sigma_j) \\ &= (i \ i+1)(j \ j+1) \\ &= (j \ j+1)(i \ i+1) \text{ since } |i-j| \geq 2 \\ &= \pi(\sigma_j \sigma_i) \end{aligned}$$

Thus $\pi(\sigma_i \sigma_j) = \pi(\sigma_j \sigma_i)$, and so π is a homomorphism. \square

Theorem 4.8.

Let X be any set. A map $f : \mathcal{B}_n \rightarrow X$ is said to be a **braid invariant** if

$$\beta \sim \beta' \Rightarrow f(\beta) = f(\beta')$$

The idea of a braid invariant is that we assign some element of X to each $\beta \in \mathcal{B}_n$ in such a way that any two equivalent braids are assigned the same element. Thus this map f is invariant under equivalence of braids. Braid invariants are a very useful tool, as they often give a quick and easy way to tell if two braids are *not* equivalent. However, we cannot in general use a braid invariant to conclude that two braids *are* equivalent.

Theorem 4.9.

Let G be any group, and $f : \mathcal{B}_n \rightarrow G$ a homomorphism. Let $\delta_n : \mathcal{B}_n \rightarrow \mathcal{B}_n$ be the map sending β to $[\beta]$. Then $f \circ \delta_n : \mathcal{B}_n \rightarrow G$ is a braid invariant.

Proof.

Let $\beta, \beta' \in \mathcal{B}_n$ with $\beta \sim \beta'$. Then $[\beta] = [\beta']$, and thus $\delta_n(\beta) = [\beta] = [\beta'] = \delta_n(\beta')$. Hence $f \circ \delta_n(\beta) = f \circ \delta_n(\beta')$, so $f \circ \delta_n$ is a braid invariant. \square

The hidden detail in the above proof is that a homomorphism is a well-defined map. We can in fact relax Theorem 4.9 so that f is merely a well-defined map on \mathcal{B}_n (instead of a homomorphism), but we have no need for it.

Corollary 4.10.

The map $\pi \circ \delta_n$, with π and δ_n as defined above, is a braid invariant for any $n \in \mathbb{N}$.

Proof.

This follows immediately from the fact that π is a homomorphism. \square

Definition 4.11.

Let $\beta \in \mathcal{B}_n$, and so we can write $\beta = \sigma_{i_1}^{\epsilon_1} \dots \sigma_{i_k}^{\epsilon_k}$ with $1 \leq i_1, \dots, i_k \leq n-1$, $\epsilon_i \in \{\pm 1\}$ for $1 \leq i \leq k$, $k \in \mathbb{N}$. Define the **exponent sum** $exp : \mathcal{B}_n \rightarrow \mathbb{Z}$ by

$$exp(\beta) = exp(\sigma_{i_1}^{\epsilon_1} \dots \sigma_{i_k}^{\epsilon_k}) := \epsilon_1 + \epsilon_1 + \dots + \epsilon_k$$

where we are viewing \mathbb{Z} as the group of additive integers.

This is a very simple map where, for a given braid β , all we do is count the number of positive crossings and subtract the number of negative crossings in any diagram for β . Equivalently, we add up the powers in a representative word for β . As it turns out, this defines a homomorphism.

Theorem 4.12.

The exponent sum is a homomorphism, and thus $\exp \circ \delta_n$ is a braid invariant for any $n \in \mathbb{N}$.

Proof.

All we need show is that \exp is well defined and a homomorphism, and the second half of the theorem follows from Theorem 4.9. But this is trivial since we have:

1. $\exp(\sigma_i \sigma_j) = 2 = \exp(\sigma_j \sigma_i)$ for $|i - j| \geq 2$
2. $\exp(\sigma_i \sigma_{i+1} \sigma_i) = 3 = \exp(\sigma_{i+1} \sigma_i \sigma_{i+1})$ for $1 \leq i \leq n - 2$ □

The exponent sum is one of the most useful braid invariants we have, as it is extremely quick and simple to calculate from either a representative word for β , or a diagram for β . We will use this invariant many times to help show some important results; its simplicity and ease of use makes it our preferred braid invariant. (Technically, it is the map $\exp \circ \delta_n$ that is the braid invariant, derived from the homomorphism \exp . But for simplicity, we will call \exp a braid invariant).

4.3 Pure braids

The homomorphism π given in Definition 4.6 has the simple geometric interpretation of taking an n -braid β and from it deriving a permutation of the numbers $1, 2, \dots, n$. Also, it is obvious that the trivial braid $\mathbf{1}_n$ lies in the kernel of π . We now turn our attention to the rest of the elements in $\text{Ker}(\pi)$, as they form an interesting subset of \mathbf{B}_n .

Definition 4.13.

An n -braid $\beta \in \mathbf{B}_n$ is said to be a **pure braid** if $\pi(\beta) = (1)$. That is, each braid string d_i starts at A_i and ends at B_i . We denote the set of pure braids in \mathbf{B}_n by \mathbf{P}_n . Thus

$$\mathbf{P}_n := \{\beta \in \mathbf{B}_n \mid \pi(\beta) = (1)\} = \text{Ker}(\pi)$$

In Figure 41 we give an example of a pure 4-braid.

Theorem 4.14.

\mathbf{P}_n is a normal subgroup of \mathbf{B}_n , and $\mathbf{B}_n/\mathbf{P}_n \cong \mathbf{S}_n$. Thus $[\mathbf{B}_n : \mathbf{P}_n] = |\mathbf{S}_n| = n!$.

Proof.

The braid permutation $\pi : \mathbf{B}_n \rightarrow \mathbf{S}_n$ is a homomorphism (see Theorem 4.7).




Figure 41: A pure 4-braid.

And $\text{Ker}(\pi) = \{\beta \in \mathbf{B}_n \mid \pi(\beta) = (1)\} = \mathbf{P}_n$. Thus \mathbf{P}_n is a normal subgroup, being the kernel of the homomorphism π .

Since $\mathbf{B}_n/\text{Ker}(\pi) \cong \text{Im}(\pi)$, and we have already defined $\text{Ker}(\pi) = \mathbf{P}_n$, all we need do is show $\text{Im}(\pi) = \mathbf{B}_n$ to complete the proof (i.e., show π is surjective).

So let $\rho \in \mathbf{S}_n$ with $\rho = \begin{pmatrix} 1 & 2 & \dots & n \\ i_1 & i_2 & \dots & i_n \end{pmatrix}$. Construct the braid $\beta \in \mathbf{B}_n$ as follows:

We have points A_1, \dots, A_n along the top face of \mathbb{D} , and points B_1, \dots, B_n along the bottom face. So join A_1 to B_{i_1} by some strictly decreasing arc d_1 in \mathbb{D} . Next, join A_2 to B_{i_2} by some other strictly decreasing arc d_2 in \mathbb{D} that does not intersect d_1 . Continue this process for all $1 \leq j \leq n$, each time joining A_j to B_{i_j} with a strictly decreasing arc d_j in \mathbb{D} that does not intersect any other d_k . Call this collection of arcs β . Then clearly $\pi(\beta) = \rho$. And since we can do this for any $\rho \in \mathbf{S}_n$, we have that π is surjective. Alternatively, we may construct a diagram representing an n -braid β' (possibly different to the β we just constructed) that satisfies $\pi(\beta') = \rho$. We do this as follows: Draw two parallel horizontal lines in the plane, labeling the top one A and the bottom B (these will eventually form the top and bottom of our diagram). Fix n distinct ordered points a_1, \dots, a_n on A and b_1, \dots, b_n on B. Then join a_n to b_{i_n} by a straight line. Now join a_{n-1} to $b_{i_{n-1}}$ by a straight line, however if this line crosses any other existing line, then make this line the understrand of the crossing. Repeat this process for a_{n-2} to $b_{i_{n-2}}$, then a_{n-3} to $b_{i_{n-3}}$, and so on. Thus we have a diagram for an n -braid β' such that each A_j is connected to B_{i_j} for all $1 \leq i, j \leq n - 1$ (i.e., $\pi(\beta') = \rho$). Observe that this method gives us a much more straightforward way to construct a diagram for such a braid. Figure 42 shows an example of the construction of such a diagram for a 5-braid β' satisfying $\pi(\beta') = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 \\ 2 & 5 & 4 & 1 & 3 \end{pmatrix}$. A quick check of the definition of a braid shows that we have in fact formed an n -braid, and from our geometric interpretation of the braid permutation π we see that $\pi(\beta) = \rho$. Thus π is surjective. \square




Figure 42: Constructing a diagram for a 5-braid β' that satisfies $\pi(\beta') = \rho$.

Theorem 4.15.

The pure braid group \mathbf{P}_n has a presentation with:

Generators: $A_{j,k}(= \sigma_{k-1}\sigma_{k-2}\dots\sigma_{j+1}\sigma_j^2\sigma_{j+1}\dots\sigma_{k-1})$ for $1 \leq j < k \leq n$

Relators:

1. $[A_{r,s}, A_{i,j}] = 1$ if $1 \leq r < s < i < j \leq n$ or $1 \leq r < i < j < s \leq n$
2. $A_{r,s}A_{r,j}A_{r,s}^{-1} = A_{s,j}^{-1}A_{r,j}A_{s,j}$ if $1 \leq r < s < j \leq n$
3. $A_{r,s}A_{s,j}A_{r,s}^{-1} = A_{s,j}^{-1}A_{r,j}^{-1}A_{s,j}A_{r,j}A_{s,j}$ if $1 \leq r < s < j \leq n$
4. $[A_{i,j}^{-1}A_{s,j}A_{i,j}, A_{r,i}]$ if $1 \leq r < s < i < j \leq n$

Proof.

The details of this proof are both long and tedious. See [4], pp 43-56 for a complete proof. \square

4.4 Quotients of \mathbf{B}_n

By allowing extra (possibly less intuitive) moves on our braids, we can find quotient groups of \mathbf{B}_n . Such a move might include replacing a double twist with two parallel arcs or vice versa (i.e., $\sigma_i^2 = \mathbf{1}_n$ for each $1 \leq i \leq n - 1$), as shown in Figure 43. We can create many quotient groups of \mathbf{B}_n in this manner.




Figure 43: Replacing a double twist with two parallel arcs.




Figure 44: Replacing a k -twist with two parallel arcs.

Definition 4.16.

Let $k \geq 2$ and $n \in \mathbb{N}$. Define the group $\mathbf{B}_n(k)$ by

$$\mathbf{B}_n(k) := \mathbf{B}_n / \langle\langle \sigma_1^k, \dots, \sigma_{n-1}^k \rangle\rangle$$

Thus we have the presentation

$$\mathbf{B}_n(k) = \langle \sigma_1, \dots, \sigma_{n-1} \mid \sigma_i \sigma_j = \sigma_j \sigma_i \text{ for } 1 \leq i, j \leq n-1 \text{ and } |i-j| \geq 2, \sigma_i \sigma_{i+1} \sigma_i = \sigma_{i+1} \sigma_i \sigma_{i+1} \text{ for } 1 \leq i \leq n-2, \sigma_i^k = 1 \text{ for } 1 \leq i \leq n-1 \rangle$$

Basically, we have added an extra relation to our braids whereby we are able to replace a k -twist with 2 parallel arcs, as is given in Figure 44.

Theorem 4.17.

For any $n, k \in \mathbb{N}$ we have that $\langle\langle \sigma_1^k, \dots, \sigma_{n-1}^k \rangle\rangle = \langle\langle \sigma_1^k \rangle\rangle$.

Proof.

Clearly $\langle\langle\sigma_1^k\rangle\rangle \subseteq \langle\langle\sigma_1^k, \dots, \sigma_{n-1}^k\rangle\rangle$.

But $\sigma_i\sigma_{i+1}\sigma_i = \sigma_{i+1}\sigma_i\sigma_{i+1}$

$$\Rightarrow \sigma_{i+1}\sigma_i\sigma_{i+1}^{-1} = \sigma_i^{-1}\sigma_{i+1}\sigma_i$$

$$\Rightarrow (\sigma_{i+1}\sigma_i\sigma_{i+1}^{-1})^k = (\sigma_i^{-1}\sigma_{i+1}\sigma_i)^k$$

$$\Rightarrow \sigma_{i+1}\sigma_i^k\sigma_{i+1}^{-1} = \sigma_i^{-1}\sigma_{i+1}^k\sigma_i$$

$$\Rightarrow \sigma_{i+1}^k = \sigma_i\sigma_{i+1}\sigma_i^k\sigma_{i+1}^{-1}\sigma_i^{-1}$$

$$\Rightarrow \sigma_{i+1}^k = (\sigma_i\sigma_{i+1})\sigma_i^k(\sigma_i\sigma_{i+1})^{-1} \in \langle\langle\sigma_i^k\rangle\rangle$$

$$\Rightarrow \sigma_{i+1}^k \in \langle\langle\sigma_i^k\rangle\rangle \text{ for } 1 \leq i \leq n-2$$

$$\Rightarrow \langle\langle\sigma_{i+1}^k\rangle\rangle \subseteq \langle\langle\sigma_i^k\rangle\rangle \text{ for } 1 \leq i \leq n-2$$

$$\Rightarrow \langle\langle\sigma_1^k, \dots, \sigma_{n-1}^k\rangle\rangle \subseteq \langle\langle\sigma_1^k\rangle\rangle$$

$$\Rightarrow \langle\langle\sigma_1^k, \dots, \sigma_{n-1}^k\rangle\rangle = \langle\langle\sigma_1^k\rangle\rangle \text{ since we already showed } \langle\langle\sigma_1^k, \dots, \sigma_{n-1}^k\rangle\rangle \supseteq \langle\langle\sigma_1^k\rangle\rangle \quad \square$$

So being able to replace a k -twist between some nominated pair of adjacent strands by two parallel arcs means we can do it for *any* pair of adjacent strands.

Corollary 4.18.

$\mathbf{B}_n(k) = \mathbf{B}_n / \langle\langle\sigma_1^k\rangle\rangle$, and thus $\mathbf{B}_n(k)$ has presentation

$$\begin{aligned} \mathbf{B}_n(k) = \langle \sigma_1, \dots, \sigma_{n-1} \mid & \sigma_i\sigma_j = \sigma_j\sigma_i \text{ for } 1 \leq i, j \leq n-1 \text{ and } |i-j| \geq 2, \\ & \sigma_i\sigma_{i+1}\sigma_i = \sigma_{i+1}\sigma_i\sigma_{i+1} \text{ for } 1 \leq i \leq n-2, \sigma_1^k = 1 \rangle \end{aligned}$$

Proof.

Follows directly from Definition 4.16 and Theorem 4.17. \square

Corollary 4.19.

For any $n \in \mathbb{N}$ we have that $\mathbf{B}_n(2) \cong \mathbf{S}_n$.

Proof.

The presentation for $\mathbf{B}_n(k)$ given above, with $k = 2$, is a standard presentation for \mathbf{S}_n . That is,

$$\mathbf{S}_n = \langle \sigma_1, \dots, \sigma_{n-1} \mid \sigma_i\sigma_j = \sigma_j\sigma_i \text{ for } 1 \leq i, j \leq n-1 \text{ and } |i-j| \geq 2, \sigma_i\sigma_{i+1}\sigma_i = \sigma_{i+1}\sigma_i\sigma_{i+1} \text{ for } 1 \leq i \leq n-2, \sigma_1^2 = 1 \rangle \quad \square$$

The next obvious question to ask is: “for which pairs (n, k) is $|\mathbf{B}_n(k)|$ finite?” The answer to this is quite interesting, and shows how different branches of mathematics can tie together. However, since the result is not used in the remainder of our discussion, we omit the proof.

Theorem 4.20.

The group $\mathbf{B}_n(k)$ is finite if and only if $k = 2$ or (n, k) is the type of one of the 5 platonic solids (That is, $(n, k) \in \{(3, 3), (3, 4), (4, 3), (3, 5), (5, 3)\}$). For these cases,

$$|\mathbf{B}_n(k)| = \left(\frac{f}{2}\right)^{n-1} n!$$

where f is the number of faces of the platonic solid of type (n, k) .

Proof.

See [4], pp 82-83.

□

5 The word and conjugacy problems on \mathbf{B}_n

The braid group is an extremely efficient way to represent braids, and has already yielded many interesting and useful results. However, we now move on to show that the word and conjugacy problems are solvable on \mathbf{B}_n , which are two very powerful facts. We actually go further and construct an algorithm for the word problem, which allows us to determine whether any given pair of n -braids are equivalent. With this we have a complete classification of braids, which aids our understanding immensely. We then proceed to find a solution to the conjugacy problem, for which we construct an explicit algorithm. This is a very useful tool when trying to classify links via the use of braids, which we investigate in section 6.

5.1 The word problem

In this section we will give a solution to the word problem on \mathbf{B}_n , by defining a (unique) normal form for elements of \mathbf{B}_n . But before we can do this, we require a few definitions and technical results. The proof we give here is purely algebraic, and is derived only from the presentation of \mathbf{B}_n . However, to aid intuition, and to motivate some of the definitions we introduce, we will appeal to the geometric interpretation behind the braid group.

Definition 5.1.

We define a **positive braid** as an element of \mathbf{B}_n that can be represented by a word involving only positive powers of the σ_i . We denote the set of all such braids by \mathbf{B}_n^+ . That is,

$$\mathbf{B}_n^+ := \{\beta \in \mathbf{B}_n \mid \beta = \sigma_{i_1} \sigma_{i_2} \dots \sigma_{i_k}, 1 \leq i_j \leq n-1, k \in \mathbb{N}\}$$

Basically, a positive braid is one which has a diagram that contains only positive crossings. If we can restrict our working to only deal with positive braids, then we may use some handy results about length and equivalence, as shall be seen shortly.

Theorem 5.2.

The standard presentation for \mathbf{B}_n also defines a semi-group, which we denote by $\bar{\mathbf{B}}_n^+$. There is also a natural embedding of $\bar{\mathbf{B}}_n^+$ into \mathbf{B}_n , sending a word over the σ_i (as an element of $\bar{\mathbf{B}}_n^+$) to the element in \mathbf{B}_n represented by the same word. Thus the image of this map is precisely \mathbf{B}_n^+ .

Proof.

We refer to [7], pp 242-243 for a proof of this. An alternative, yet somewhat less suitable proof (since it relies on a solution to the word problem in \mathbf{B}_n) can be found in [6], pp 194-195. \square

A consequence of this is that two positive words w_1, w_2 are equivalent in \mathbf{B}_n precisely when they are equivalent in $\bar{\mathbf{B}}_n^+$, which we shall use later.

Theorem 5.3.

If two positive n -braids β, β' are equivalent, then when written as positive words they are of the same length.

Proof.

Suppose $\beta = \sigma_{i_1} \sigma_{i_2} \dots \sigma_{i_k}$, $\beta' = \sigma_{j_1} \sigma_{j_2} \dots \sigma_{j_l}$ as words over the σ_i , and that $\beta \sim \beta'$. Then $\text{exp}(\beta) = k$ and $\text{exp}(\beta') = l$. But $\beta \sim \beta'$, so $\text{exp}(\beta) = \text{exp}(\beta')$ and hence $k = l$. So β and β' are of the same length. \square

Corollary 5.4.

Given a positive braid $\beta \in \mathbf{B}_n$, there are only finitely many other positive words over the σ_i that define braids equivalent to β .

Proof.

Suppose a positive word for β has length k . Then all equivalent positive braids also have length k , when written as positive words. Since the number of σ_i 's is finite, there can be only finitely many positive words of a fixed length. \square

A direct consequence of Theorem 5.2 and Corollary 5.4 is that, given a positive n -braid β (written as a positive word w), we can compute the finite set of all positive words that represent β in \mathbf{B}_n . This can be done as follows: Take the representative (positive) word w for β , and view it as an element of $\bar{\mathbf{B}}_n^+$. We know there can only be finitely many words equivalent to w in $\bar{\mathbf{B}}_n^+$ from the comment made after Theorem 5.2 and Corollary 5.4. So take w and apply all possible defining relations to get a larger set of words. We then apply all possible defining relations to this set to get a larger set of words, all still equivalent to w in $\bar{\mathbf{B}}_n^+$. We continue this until the set closes to give all words in $\bar{\mathbf{B}}_n^+$ equivalent to w (we know the set will eventually close because it is finite). Thus this is also precisely the set of positive words equivalent to w in \mathbf{B}_n , again by our comment after Theorem 5.2.

We now re-introduce the braid Δ_n , previously mentioned in Theorem 4.1. We shall soon see that this braid plays a vital role in our solution to the word problem.




Figure 45: *The Garside braid Δ_5 .*

Definition 5.5.

Let Π_r denote the braid $\sigma_1\sigma_2\dots\sigma_r$. We define the **Garside braid** $\Delta_n \in \mathbf{B}_n$ by

$$\Delta_n := (\sigma_1\sigma_2\dots\sigma_{n-1})(\sigma_1\sigma_2\dots\sigma_{n-2})\dots(\sigma_1\sigma_2)(\sigma_1) = \Pi_{n-1}\Pi_{n-2}\dots\Pi_2\Pi_1$$

See Figure 45 for a diagram for Δ_5 .

We have already given a brief introduction to the Garside braid in the discussion following Theorem 4.1. The only comment we wish to make here is a reminder of the construction of Δ_n , by taking $\mathbf{1}_n$ and placing a half-twist at the bottom (anti-clockwise, when viewed from the bottom of the braid up).

Lemma 5.6.

In \mathbf{B}_n , for any i, r with $1 < i \leq r \leq n - 1$, we have that $\sigma_i\Pi_r = \Pi_r\sigma_{i-1}$.

Proof.

$$\begin{aligned} \sigma_i\Pi_r &= \sigma_i(\sigma_1\dots\sigma_r) \\ &= \sigma_i(\sigma_1\dots\sigma_{i-2})\sigma_{i-1}\sigma_i(\sigma_{i+1}\sigma_{i+2}\dots\sigma_r) \\ &= (\sigma_1\dots\sigma_{i-2})\sigma_i\sigma_{i-1}\sigma_i(\sigma_{i+1}\sigma_{i+2}\dots\sigma_r) \\ &= (\sigma_1\dots\sigma_{i-2})\sigma_{i-1}\sigma_i\sigma_{i-1}(\sigma_{i+1}\sigma_{i+2}\dots\sigma_r) \text{ (can do this since } 1 < i \leq r) \\ &= (\sigma_1\dots\sigma_{i-2})\sigma_{i-1}\sigma_i(\sigma_{i+1}\dots\sigma_r)\sigma_{i-1} \\ &= \Pi_r\sigma_{i-1} \end{aligned}$$

□

Lemma 5.7.

For any $1 \leq t \leq n - 1$, we have that $\sigma_1(\Pi_t \Pi_{t-1} \dots \Pi_1) = (\Pi_t \Pi_{t-1} \dots \Pi_1) \sigma_t$.

Proof.

First let $t = 1$. Then $\sigma_1 \Pi_1 = \sigma_1 \sigma_1 = \Pi_1 \sigma_1$. Now let $2 \leq t \leq n - 1$. Then

$$\begin{aligned}
\sigma_1 \Pi_t \Pi_{t-1} \dots \Pi_1 &= \sigma_1 \Pi_t (\sigma_1 \sigma_2 \dots \sigma_{t-1}) \Pi_{t-2} \dots \Pi_1 \\
&= \sigma_1 (\sigma_2 \sigma_3 \dots \sigma_t) \Pi_t \Pi_{t-2} \dots \Pi_1 \text{ (by Lemma 5.6)} \\
&= \Pi_t \Pi_t \Pi_{t-2} \dots \Pi_1 \\
&= \Pi_t (\sigma_1 \sigma_2 \sigma_3 \dots \sigma_t) \Pi_{t-2} \dots \Pi_1 \\
&= \Pi_t (\sigma_1 \sigma_2 \sigma_3 \dots \sigma_{t-1}) \sigma_t \Pi_{t-2} \dots \Pi_1 \\
&= \Pi_t \Pi_{t-1} \sigma_t \Pi_{t-2} \dots \Pi_1 \\
&= \Pi_t \Pi_{t-1} \Pi_{t-2} \dots \Pi_1 \sigma_t \text{ (since } \sigma_{t+1}^{\pm 1}, \sigma_{t-1}^{\pm 1} \text{ don't appear in } \Pi_{t-2} \dots \Pi_1)
\end{aligned}$$

□

Corollary 5.8.

For any $n \in \mathbb{N}$, we have that $\sigma_1 \Delta_n = \Delta_n \sigma_{n-1}$.

Proof.

Set $t = n - 1$ in Lemma 5.7

□

Theorem 5.9.

For any $1 \leq i \leq n - 1$, we have that $\sigma_i \Delta_n = \Delta_n \sigma_{n-i}$.

Proof.

If $i = 1$, then the result is immediate from Corollary 5.8. So suppose $2 \leq i \leq n - 1$. Then we have

$$\begin{aligned}
\sigma_i \Delta_n &= \sigma_i \Pi_{n-1} \Pi_{n-2} \dots \Pi_1 \\
&= \Pi_{n-1} \sigma_{i-1} \Pi_{n-2} \dots \Pi_1 \text{ (by Lemma 5.6)} \\
&\quad \vdots \\
&= \Pi_{n-1} \dots \Pi_{n-i+1} \sigma_1 (\Pi_{n-i} \dots \Pi_1) \text{ (by Lemma 5.6)} \\
&= \Pi_{n-1} \dots \Pi_{n-i+1} (\Pi_{n-i} \Pi_{n-i-1} \dots \Pi_1) \sigma_{n-i} \text{ (by Lemma 5.7)} \\
&= \Delta_n \sigma_{n-i}
\end{aligned}$$

□

Theorem 5.10.

The square of the Garside braid lies in the centre of \mathbf{B}_n . That is, $\Delta_n^2 \in Z(\mathbf{B}_n)$.

Proof.

For any $1 \leq i \leq n - 1$ we have

$$\begin{aligned}
\sigma_i \Delta_n^2 &= (\sigma_i \Delta_n) \Delta_n \\
&= (\Delta_n \sigma_{n-i}) \Delta_n \text{ (by Theorem 5.9)} \\
&= \Delta_n (\sigma_{n-i} \Delta_n) \\
&= \Delta_n (\Delta_n \sigma_{n-(n-i)}) \text{ (by Theorem 5.9)} \\
&= \Delta_n^2 \sigma_i
\end{aligned}$$

□

Theorem 5.11.

The inner automorphism defined by $\tau : \mathbf{B}_n \rightarrow \mathbf{B}_n$, $\tau(\beta) = \Delta_n^{-1} \beta \Delta_n$ sends σ_i to σ_{n-i} for each $1 \leq i \leq n - 1$.

Proof.

To show $\tau(\sigma_i) = \sigma_{n-i}$, we merely appeal to Theorem 5.9 to give

$$\begin{aligned}
\tau(\sigma_i) &= \Delta_n^{-1} \sigma_i \Delta_n \\
&= \Delta_n^{-1} (\Delta_n \sigma_{n-i}) \text{ (by Theorem 5.9)} \\
&= \sigma_{n-i}
\end{aligned}$$

□

Lemma 5.12.

Let $\beta \in \mathbf{B}_n$. Then $\beta \Delta_n^m = \Delta_n^m \tau(\beta)$ for any odd $m \in \mathbb{Z}$. Also, if β is a positive braid, then so is $\tau(\beta)$.

Proof.

For the first part of the lemma, we need only prove $\beta \Delta_n = \Delta_n \tau(\beta)$, since $\Delta_n^{2k} \in Z(\mathbf{B}_n)$ for any $k \in \mathbb{Z}$. But this is immediately obvious from the definition of τ as given in Theorem 5.11.

For the second part of the lemma, we simply note that τ sends σ_i to σ_{n-i} for each $1 \leq i \leq n - 1$ (see Theorem 5.11). Thus, for any positive braid $\beta = \sigma_{i_1} \sigma_{i_2} \dots \sigma_{i_k}$, we have

$$\tau(\beta) = \tau(\sigma_{i_1} \sigma_{i_2} \dots \sigma_{i_k}) = \sigma_{n-i_1} \sigma_{n-i_2} \dots \sigma_{n-i_k}$$

which is again a positive braid. □

Definition 5.13.

Let β be an n -braid, and $w = \sigma_{i_1}^{\epsilon_1} \dots \sigma_{i_k}^{\epsilon_k}$ be a word over the Artin generators representing β . We define the **reverse** of the word w , denoted $rev(w)$, by

$$rev(w) := \sigma_{i_k}^{\epsilon_k} \dots \sigma_{i_1}^{\epsilon_1}$$

We define the **reverse** of the braid β , denoted $rev(\beta)$, by the element in \mathbf{B}_n represented by the word $rev(w)$.

The reverse of a word is simply the word written backwards. The reverse of a braid β can be seen in a similar fashion. That is, if we imagine our braid β being constructed from the trivial braid by adding successive twists between adjacent strands, then $rev(\beta)$ is the braid formed by performing the last twist first, the second last twist second, and so on.

Lemma 5.14.

For any two words w, w' over the Artin generators, we have that $rev(ww') = rev(w')rev(w)$.

Proof.

Suppose $w = \sigma_{i_1}^{\epsilon_1} \dots \sigma_{i_k}^{\epsilon_k}$, $w' = \sigma_{j_1}^{\epsilon'_1} \dots \sigma_{j_l}^{\epsilon'_l}$ as words. Then

$$rev(ww') = rev(\sigma_{i_1}^{\epsilon_1} \dots \sigma_{i_k}^{\epsilon_k} \sigma_{j_1}^{\epsilon'_1} \dots \sigma_{j_l}^{\epsilon'_l}) = \sigma_{j_l}^{\epsilon'_l} \dots \sigma_{j_1}^{\epsilon'_1} \sigma_{i_k}^{\epsilon_k} \dots \sigma_{i_1}^{\epsilon_1} = rev(w')rev(w). \quad \square$$

Corollary 5.15.

For any two n -braids β, β' , we have that $rev(\beta\beta') = rev(\beta')rev(\beta)$.

Proof.

Let w, w' be words representing n -braids β, β' respectively. Then

$$rev(\beta\beta') =_{\mathbf{B}_n} rev(ww') =_{\mathbf{B}_n} rev(w')rev(w) =_{\mathbf{B}_n} rev(\beta')rev(\beta). \quad \square$$

Lemma 5.16.

For any $1 \leq r \leq n - 1$ we have that $rev(\Pi_r \dots \Pi_1) = \Pi_r \dots \Pi_1$.

Proof.

We proceed by induction. Now clearly $rev(\Pi_1) = rev(\sigma_1) = \sigma_1 = \Pi_1$. So suppose $rev(\Pi_t \dots \Pi_1) = \Pi_t \dots \Pi_1$ for some fixed t with $1 \leq t \leq n - 2$.

Then

$$\begin{aligned}
rev(\Pi_{t+1}\Pi_t\dots\Pi_1) &= rev(\Pi_t\dots\Pi_1)rev(\Pi_{t+1}) \\
&= (\Pi_t\dots\Pi_1)rev(\Pi_{t+1}) \text{ (by induction hypothesis)} \\
&= (\Pi_t\dots\Pi_1)rev(\sigma_1\dots\sigma_{t+1}) \\
&= (\Pi_t\dots\Pi_1)(\sigma_{t+1}\dots\sigma_1) \\
&= \Pi_t\sigma_{t+1}\Pi_{t-1}\Pi_{t-2}\dots\Pi_1(\sigma_t\dots\sigma_1) \text{ (} \sigma_{t+1}, \Pi_{t-1}\Pi_{t-2}\dots\Pi_1 \text{ commute)} \\
&= \Pi_t\sigma_{t+1}\Pi_{t-1}\sigma_t\Pi_{t-2}\Pi_{t-3}\dots\Pi_1(\sigma_{t-1}\dots\sigma_1) \\
&\quad \vdots \\
&= (\Pi_t\sigma_{t+1})(\Pi_{t-1}\sigma_t)(\Pi_{t-2}\sigma_{t-1})\dots(\Pi_1\sigma_2)\sigma_1 \\
&= \Pi_{t+1}\Pi_t\dots\Pi_2\Pi_1
\end{aligned}$$

Thus, since our statement is true for $t = 1$, then by induction it is true for all $1 \leq t \leq n - 2$ (noting that if $t > n - 2$, Π_{t+1} is undefined). \square

Corollary 5.17.

For any $n \in \mathbb{N}$ we have that $rev(\Delta_n) = \Delta_n$.

Proof.

Put $r = n - 1$ in Lemma 5.16. \square

Theorem 5.18.

For each $1 \leq i \leq n - 1$ there exist (non-unique) positive braids L_i, R_i such that $\Delta_n = L_i\sigma_i = \sigma_i R_i$.

Proof.

Since $\Delta_n = \Pi_{n-1}\dots\Pi_1$, we already have $\Delta_n = L_1\sigma_1$, with $L_1 = \Pi_{n-1}\dots\Pi_2$. Now, let $f(\sigma_2, \sigma_3, \dots, \sigma_t)$ be a positive word over the Artin generators $\sigma_2, \sigma_3, \dots, \sigma_t$ only, for some $2 \leq t \leq n - 1$. Then, by Lemma 5.6, we have

$\Pi_t f(\sigma_1, \sigma_2, \dots, \sigma_{t-1}) = f(\sigma_2, \sigma_3, \dots, \sigma_t)\Pi_t$. For any $2 \leq i \leq n - 1$, denote $\Pi_{i-1}\Pi_{i-2}\dots\Pi_1$ by $f(\sigma_1, \sigma_2, \dots, \sigma_{i-1})$. Then

$$\begin{aligned}
\Delta_n &= \Pi_{n-1}\Pi_{n-2}\dots\Pi_{i+1}\Pi_i(\Pi_{i-1}\dots\Pi_1) \\
&= \Pi_{n-1}\Pi_{n-2}\dots\Pi_{i+1}\Pi_i f(\sigma_1, \sigma_2, \dots, \sigma_{i-1}) \\
&= \Pi_{n-1}\Pi_{n-2}\dots\Pi_{i+1} f(\sigma_2, \sigma_3, \dots, \sigma_i)\Pi_i \text{ (by Lemma 5.6)} \\
&= \Pi_{n-1}\Pi_{n-2}\dots\Pi_{i+1} f(\sigma_2, \sigma_3, \dots, \sigma_i)\sigma_1\sigma_2\dots\sigma_{i-1}\sigma_i \\
&= L_i\sigma_i
\end{aligned}$$

where $L_i = \Pi_{n-1}\Pi_{n-2}\dots\Pi_{i+1} f(\sigma_2, \sigma_3, \dots, \sigma_i)\sigma_1\sigma_2\dots\sigma_{i-1}$ is a positive braid. So our positive braids L_i exist for all $1 \leq i \leq n - 1$.

Now, given some $1 \leq i \leq n - 1$, let $R_i = \text{rev}(L_i)$, which is also a positive braid. Then

$$\begin{aligned}
\sigma_i R_i &= \sigma_i \text{rev}(L_i) \\
&= \text{rev}(\sigma_i) \text{rev}(L_i) \\
&= \text{rev}(L_i \sigma_i) \text{ (by Corollary 5.15)} \\
&= \text{rev}(\Delta_n) \\
&= \Delta_n \text{ (by Lemma 5.17)}
\end{aligned}$$

So our positive braids R_i exist for all $1 \leq i \leq n - 1$. □

Theorem 5.19.

Any $\beta \in \mathbf{B}_n$ can be written as $\beta = \Delta_n^r T$, where T is a positive braid and $r \in \mathbb{Z}$.

As an aside, the above expression is in general non-unique.

Proof.

Since $\beta \in \mathbf{B}_n$, we can write $\beta = \sigma_{i_1}^{\epsilon_1} \sigma_{i_2}^{\epsilon_2} \dots \sigma_{i_k}^{\epsilon_k}$ for some $1 \leq i_j \leq n - 1$, $k \in \mathbb{N}$. But for each $1 \leq i \leq n - 1$ we have $\Delta_n = L_i \sigma_i$ for some positive braid L_i , and thus $\sigma_i^{-1} = \Delta_n^{-1} L_i$. So we can replace each σ_i^{-1} with $\Delta_n^{-1} L_i$ in our expression for β . We then shift all powers of Δ_n^{-1} to the left using Lemma 5.12 and Lemma 5.10, leaving us with $\Delta_n^r T$, where T is a positive braid and $r \in \mathbb{Z}$. □

This theorem is extremely useful, as it allows us to package all negative powers of the Artin generators into the prefix Δ_n^r . Thus we are left to deal with the positive braid T , and we have already proven many useful theorems about positive braids. We now require only a few more small results to give our normal form for a braid β .

Theorem 5.20.

Let $\beta \in \mathbf{B}_n$. Then there exists some maximal value of r , say l , such that if $r > l$, then there is no positive braid T' satisfying $\beta = \Delta_n^r T'$. Also, this l is a braid invariant.

Proof.

Since $\beta = \Delta_n^r T$, then all we need do is look at exponent sums. So $\text{exp}(\beta) = \text{exp}(\Delta_n^r(T)) = r \cdot \text{exp}(\Delta_n) + \text{exp}(T)$ since exp is a homomorphism.

But

$$\begin{aligned}
\exp(\Delta_n) &= \exp(\Pi_{n-1} \dots \Pi_1) \\
&= \exp((\sigma_1 \sigma_2 \dots \sigma_{n-1})(\sigma_1 \sigma_2 \dots \sigma_{n-2}) \dots (\sigma_1 \sigma_2)(\sigma_1)) \\
&= r \cdot ((n-1) + (n-2) + \dots + (1)) \\
&= \frac{rn(n-1)}{2}
\end{aligned}$$

Thus $\exp(\beta) = \frac{rn(n-1)}{2} + \exp(T) \geq \frac{rn(n-1)}{2}$ since T is a positive braid (and so $\exp(T) \geq 0$). Hence $\exp(\beta) \geq \frac{rn(n-1)}{2}$, or alternatively $r \leq \frac{2\exp(\beta)}{n(n-1)}$ (where here we ignore the $n = 1$ case since \mathbf{B}_1 is the trivial group). Thus $\frac{2\exp(\beta)}{n(n-1)}$ is an upper bound for r , and since any set of integers bounded above attains its maximum, r has some maximum value l . This l is a braid invariant since if $\beta \sim \beta'$, then any word $\Delta_n^r(T)$ representing β in \mathbf{B}_n also represents β' in \mathbf{B}_n (and vice-versa). \square

Definition 5.21.

We denote the braid invariant l for β as given above in Theorem 5.20 as the **infimum** of β , written $\text{inf}(\beta)$.

Definition 5.22.

Let $w = \sigma_{i_1} \dots \sigma_{i_k}$ be a positive word over the Artin generators $\{\sigma_1, \dots, \sigma_{n-1}\}$. We define the **subscript array** of w , denoted $s(w)$, as the k -tuple (i_1, i_2, \dots, i_k) .

We can now give a (unique) normal form for braids in \mathbf{B}_n .

Definition 5.23.

Let $\beta \in \mathbf{B}_n$ and $i = \text{inf}(\beta)$ (thus we can write $\beta = \Delta_n^i T$, for some positive braid T). We then find the finite set $\{Z_0, Z_1, \dots, Z_k\}$ of all positive words over the Artin generators equivalent to T in \mathbf{B}_n , as shown in our comment after Corollary 5.4. Now choose the Z_j whose subscript array is lexicographically minimal (this is well defined and unique). Then our normal form for β is the word $\Delta_n^i Z_j$.

Theorem 5.24.

The word problem is solvable on the braid group \mathbf{B}_n for any $n \in \mathbb{N}$.

Proof.

Suppose we are given $\beta, \beta' \in \mathbf{B}_n$, represented as words over the σ_i and their inverses. Find the normal forms for β, β' . Then $\beta = \beta'$ if and only if they have the same normal form. \square

So we have found a normal form for elements of \mathbf{B}_n , and thus shown that the word problem is solvable on \mathbf{B}_n . However, given $\beta \in \mathbf{B}_n$ and a word w over

the Artin generators representing β , the time taken to check if $\beta = \mathbf{1}_n$ by our method is exponential in both n and $|w|$. So our solution to the word problem is somewhat slow. There are better (i.e., faster) methods of solving the word problem on \mathbf{B}_n . In [2], pp 63-64 we see a much faster solution to the word problem, that uses the left-greedy normal form on \mathbf{B}_n . It is known as the Adyan-Thurston-ElRifai-Morton solution to the word problem, and (to show $[w] = \mathbf{1}_n$) is $\mathcal{O}(|w|^2 n \log n)$.

5.2 The conjugacy problem

The solution to the conjugacy problem on \mathbf{B}_n is much longer and more difficult than that of the word problem. Though we will describe an explicit algorithm to determine if two braids are conjugate, we will only sketch the proof behind it. This sketch is based largely on the sketch provided in [2], pp 61-66.

Theorem 5.25.

The conjugacy problem is solvable on the braid group \mathbf{B}_n for any $n \in \mathbb{N}$.

Proof.

Let l, r be positive words over the Artin generators, such that $\Delta_n = lr$ as n -braids. We call l a **left divisor** of Δ_n , and r a **right divisor** of Δ_n . Now let \mathcal{P} and \mathcal{P}' denote the set of all left (respectively right) divisors of Δ_n . Then it can be shown that $\mathcal{P} = \mathcal{P}'$. Also, \mathcal{P} contains $n!$ braids, which are in 1-1 correspondence with the permutations of their end-points. This correspondence is given by the map sending σ_i to the permutation $(i \ i+1) \in \mathbf{S}_n$, and for this reason we call \mathcal{P} the set of **permutation braids**. Combining this with the fact that for any $p \in \mathcal{P}$ all the crossings of p are positive, we can reconstruct p from just its permutation (i.e., from its image under the braid permutation π). We obtain a unique braid from this, even though its representation as a word is non-unique. Permutation braids have the following property: If we ‘tighten’ the braids strings of any $p \in \mathcal{P}$, we obtain a braid where no two strings cross more than once.

We now describe a unique normal form for elements of \mathbf{B}_n . Recall that, for any $\beta \in \mathbf{B}_n$, we could write $\beta = \Delta_n^i P$, with P a positive braid and $i = \text{inf}(\beta)$. Now, if the strands in P cross at most once, then $P \in \mathcal{P}$, and we stop. If not, set $P = l_1 l'_1$, where l_1 is a positive braid of maximal length in which no two strands cross more than once, and l'_1 is the rest of P . If no two strands in l'_1 cross more than once, set $l_2 = l'_1$ and stop. If not, repeat, setting $l'_1 = l_2 l'_2$, where l_2 is a positive braid of maximal length in which no two strands cross more than once, and l'_2 is the rest of l'_1 . Continue this process until we obtain $P = l_1 \dots l_s$, where each $l_i \in \mathcal{P}$ and each l_i has maximal length whenever we factorise $l_{i-1}^{-1} \dots l_1^{-1} P$ as

$x_i x'_i$ with $x_i \in \mathcal{P}$ and x'_i positive. This representation is unique, up to the choice of words representing l_1, \dots, l_s , and each of the l_i is determined uniquely by the permutation of its strands. The factorisation $\beta = \Delta_n^i l_1 \dots l_s$ is known as the **left greedy normal form** (or LGNF for short). This is a unique normal form, and thus solves the word problem. It is known as the Adyan-Thurston-ElRifai-Morton solution to the word problem (as mentioned after the proof of Theorem 5.24), and is $\mathcal{O}(|w|^2 n \log n)$.

We now use our unique normal form as a tool to solve the conjugacy problem. Denote the conjugacy class of a word w over the Artin generators by $\{w\}$. The exponent sum is also an invariant of conjugacy class (as well as equivalence class). Thus there are only finitely many words w' in $\{w\}$ satisfying $\text{inf}(w') \geq \text{inf}(w)$. This is because $w' = \Delta_n^{\text{inf}(w')} l_1 \dots l_s$ in LGNF, and so $\text{exp}(w') = \text{inf}(w') \cdot \text{exp}(\Delta_n) + \text{exp}(l_1 \dots l_s) \geq \text{inf}(w') \cdot \text{exp}(\Delta_n)$ since all the l_i are positive braids (and so $\text{exp}(l_1 \dots l_s) \geq 0$). Thus $\text{inf}(w') \leq \frac{\text{exp}(w')}{\text{exp}(\Delta_n)}$. So we have $\text{inf}(w) \leq \text{inf}(w') \leq \frac{\text{exp}(w')}{\text{exp}(\Delta_n)}$, and as we increase $\text{inf}(w')$ we must decrease the length of $l_1 \dots l_s$ (since exp is also a braid invariant). And we have already shown from Theorem 5.4 that there are only a finite number of positive braids of a given length, thus proving our statement.

Let $\text{Inf}(w)$ denote the maximum value of $\text{inf}(w')$ for all $w' \in \{w\}$. Suppose we have some representative w_1 of $\{w\}$ with LGNF $w_1 = \Delta_n^I L_1 \dots L_S$, where $I = \text{Inf}(w)$ and S is minimal for all braids in LGNF that are conjugate to w . Let $\text{Sup}(w)$ denote the integer $I + S$. Then clearly $\text{Inf}(w)$ and $\text{Sup}(w)$ are invariant in $\{w\}$. We define the ElRifai-Morton **super summit set**, denoted S_w , as the set of all elements of $\{w\}$ that realise both $\text{Inf}(w)$ and $\text{Sup}(w)$.

We now illustrate a constructive method for finding $\text{Inf}(w)$ and $\text{Sup}(w)$. Firstly, define the **cycling** of a braid $w = \Delta_n^i l_1 \dots l_k$ (in LGNF) by $c(w) = \tau^{-1}(l_1) w (l_1)$. Similarly, define the **decycling** of w by $d(w) = l_k w l_k^{-1}$. Notice that $c(w)$ and $d(w)$ also lie in $\{w\}$. If we then put $c(w)$ and $d(w)$ into LGNF, we have that $\text{inf}(c(w)), \text{inf}(d(w)) \geq i$. And thus, if we have increased i , we must have subsequently decreased s . It is shown in [9] that if $\text{inf}(w)$ is not maximal for $\{w\}$, it can be increased by repeated cyclings. Similarly, if $\text{sup}(w)$ is not minimal for $\{w\}$, it can be decreased by repeated decyclings. In [10] it is shown that if $\text{inf}(w)$ is not maximal (respectively $\text{sup}(w)$ not minimal), then it can be increased (respectively decreased) by performing no more than $(n)(n-1)/2$ cyclings (respectively decyclings). Thus we have a way to increase $\text{inf}(w)$ and decrease $\text{sup}(w)$, and also verify when no more such increases or decreases are possible. This method also yields an element that realises both $\text{Inf}(w)$ and $\text{Sup}(w)$ when put in LGNF.

So we can use the above to obtain the super summit set S_w . It is shown in [11]

that we need only consider the subset of S_w of elements that lie in a closed orbit under cycling (it turns out we need not consider decycling). We call this finite set the **ultra summit set**, denoted U_w . We also see in [11] that $\{w\} = \{w'\}$ if and only if $U_w = U_{w'}$, which will ultimately be how we solve the conjugacy problem.

We now move to computing U_w . In [11] we see that if $w_i, w_j \in U_w$ then there exists a finite chain $w_i = w_{i,1} \rightarrow w_{i,2} \rightarrow \dots \rightarrow w_{i,q} = w_j$ such that each $w_{i,k}$ lies in U_w and is obtained from $w_{i,k-1}$ by conjugating by an element in \mathcal{P} . Thus we now have an explicit way to find U_w given one element $\rho \in U_w$ (which we know how to find from w). That is, we compute all conjugates of ρ by the $n!$ elements in \mathcal{P} . We then put each into LGNF and retain those that realise $Inf(w)$ and $Sup(w)$, and lie in a closed orbit under cycling, and are a new elements of U_w (where it is best to check these in the order given). We then repeat for each element so added to our set, and eventually our set closes to give U_w . Thus to check if $\{w\} = \{w'\}$, all we need do is find the finite sets $U_w, U_{w'}$ and compare them, recalling that $\{w\} = \{w'\}$ if and only if $U_w, U_{w'}$. \square

In regards to computational time, the above method for computing the ultra summit set for an n -braid β (represented by a word w) is exponential in both n and $|w|$. However, more study into this set may lead to a more efficient way to calculate it, though at present this has not been done. See [2], p 72 for further discussion of this.

6 Braids and links

We now illustrate some important relationships between braids and links, for it was the study of knots and links that first motivated a study of braids. In some special circumstances, much information can be gathered about knots by looking at associated braids (we leave the detail until later). But alas, a general method to classify knots and links has not been found from braid theory, though a theorem by Markov brings us tantalisingly close to such a classification (see Theorem 6.9).

For the remainder of this section we assume the reader is familiar with some elementary knot theory. We shall make many references to knot/link diagrams, and assume the standard definition (not to be confused with a braid diagram). See [8] for a general introduction to knot theory.

6.1 Braid closure

The first thing we must do is develop some way to turn a braid into a link (or vice-versa). Since the two are so similar, it seems logical to find such a connection between them. As it turns out, we can define ways to do both. That is, turn a link into a braid, and turn a braid into a link. However, one direction is significantly easier than the other, as we shall see shortly.

Firstly though, we remind the reader of the definitions of the points A_i and B_i in \mathbb{D} (from Definition 2.2).

$$A_i := \left(\frac{1}{2}, \frac{i}{n+1}, 1\right) \text{ for } 1 \leq i \leq n$$

$$B_i := \left(\frac{1}{2}, \frac{i}{n+1}, 0\right) \text{ for } 1 \leq i \leq n$$

Definition 6.1.

Let $\beta \in \mathbf{B}_n$. Define the **closure** of β as follows: Join A_i to B_i by an arc c_i consisting of three straight line segments:

1. A straight line $c_{i,1}$ from A_i to $(0, \frac{i}{n+1}, 1)$ (running along the top face of \mathbb{D}).
2. A straight line $c_{i,2}$ from $(0, \frac{i}{n+1}, 1)$ to $(0, \frac{i}{n+1}, 0)$ (running along the back face of \mathbb{D}).
3. A straight line $c_{i,3}$ from $(0, \frac{i}{n+1}, 0)$ to $(\frac{1}{2}, \frac{i}{n+1}, 0)$ (running along the bottom face of \mathbb{D}).




Figure 46: The 4-braid $\beta = \sigma_2^{-1}\sigma_1^2\sigma_3^2$, and its closure.




Figure 47: A diagram for the 5-braid $\beta = \sigma_1\sigma_3^{-1}\sigma_2^2\sigma_4^2\sigma_1^{-1}\sigma_4$, and its closure.

We denote the closure of β by $\tilde{\beta}$, and in Figure 46 give an example of the 4-braid $\beta = \sigma_2^{-1}\sigma_1^2\sigma_3^2$ and its closure.

An easier way to see this is by looking at a diagram for β . A knot-diagram (see [8], p 3) of $\tilde{\beta}$ can be found by taking a diagram for β (as a braid) and joining each A_i to B_i by an arc that goes around to the right, doing this for each $i = 1, 2, \dots, n$ in decreasing order starting from $i = n$. See Figure 47 which shows a diagram for the 5-braid $\beta = \sigma_1\sigma^{-1}\sigma_2^2\sigma_4^2\sigma_1^{-1}\sigma_4$, and its closure.

Theorem 6.2.

The closure $\tilde{\beta}$ of any $\beta \in \mathbf{B}_n$ is a link with at most n components.

Proof.

Closing a braid gives us a collection of polygonal curves in \mathbb{R}^3 . These curves have no self-intersections, nor do they intersect each other. And all these curves are closed, so the closure $\tilde{\beta}$ is a link. Clearly $\tilde{\beta}$ can have no more than n components, since each d_i connects with some unique c_j , so we have at most n disconnected curves in the link $\tilde{\beta}$, and thus at most n components. \square

So we now have a way of turning a braid into a link. In retrospect, this was the most obvious way to do it, as the only problem we had to fix in our braid to call it a link was to close all the open-ended curves. However, it is quite common to have an orientation on a link (and thus call it an oriented link). So we now need a consistent way to orient our new braid closure. It turns out that this isn't too difficult.

Definition 6.3.

Let $\beta \in \mathbf{B}_n$. We define an orientation on β as follows: For each braid string d_i , orient d_i from A_i to $B_{\pi(d_i)}$ (i.e., orient each strand from top to bottom). This induces an orientation on $\tilde{\beta}$ if we orient each c_i from B_i to A_i (i.e., from bottom to top).

Theorem 6.4.

Let $\beta \in \mathbf{B}_n$. Then the orientation on the link $\tilde{\beta}$ as described above is self-consistent.

Proof.

This is immediately obvious as arcs inside \mathbb{D} are oriented from the top face to the bottom face, and arcs outside \mathbb{D} are oriented from the bottom face to the top face. \square

Thus we have formulated an intuitive, consistent, and simple way to turn any braid into an oriented link. However, the original motivation behind studying braids was to learn more about links. So really, we are more interested in forming some sort of braid from a link than a link from a braid. In the next section we describe a way to do this.




Figure 48: An example of an oriented link diagram with all labels.

6.2 Alexander's theorem

Now that we have shown any braid can be closed to form an (oriented) link, we ask the reverse question. That is, given an oriented link L , can we find some braid β for which $\tilde{\beta}$ is equivalent to L ? It turns out that we can, but to find such a braid is a much more difficult procedure. We in fact require a theorem by Alexander to do this (the proof is constructive and describes an algorithm for finding such a braid). However, the algorithm can be quite long and difficult, even for links with few crossings.

Theorem 6.5. (*Alexander's Theorem*)

Given an oriented link L , there exists an $n \in \mathbb{N}$ and an n -braid $\beta \in \mathbf{B}_n$ such that $L \approx_l \tilde{\beta}$, where \approx_l denotes equivalence of links.

Proof.

We give a constructive proof, that can be applied to any oriented link L . We begin by taking an oriented diagram of our link L , and we shall call this diagram D . By definition, any diagram of L has only a finite number of intersection points. We then perturb this diagram so that it has only a finite number of local minima and local maxima (where 'height' is the distance up the vertical axis). We now draw a square S around (but not intersecting) D , with top edge labeled A and bottom edge labeled B . See Figure 48 for an example of all this. Clearly this diagram D must have at least one local minimum and at least one local maximum, since L consists of closed curves. So choose one of these minima, say a . Then follow the orientation of L from a until a local maximum is reached, say b (we know we will eventually reach a maximum because L consists of closed curves). We say the arc \overline{ab} is **increasing** from a to b and call it an **increasing arc**. We now mark $k + 1$ points on \overline{ab} , $a = a_0, a_1, \dots, a_k = b$, such that each arc




Figure 49: *An increasing arc.*

$\overline{a_i a_{i+1}}$ crosses only one other part of D . (i.e., for each $1 \leq i \leq k-1$, a_i is placed on \overline{ab} between two crossings of \overline{ab}). See Figure 49 for a labeled increasing arc.

We now define two procedures that can be applied to an increasing arc \overline{ab} .

Procedure I.

If $k = 0$ for our increasing arc \overline{ab} , then choose a point on \overline{ab} somewhere strictly between a and b and label it P . Now cut \overline{ab} at P and pull the two loose ends apart and over all other arcs in D so that the endpoint coming from b lies on A and the other lies on B . (We must be careful here to make sure that we still have a regular diagram. That is, there is still only a finite number of crossing points). Label the point on A by P' , and the point on B by P'' . Then join P'' to P' by an oriented arc lying outside S , that begins at P'' , travels around and to the right of S , then joins P' from above. Figure 50 illustrates this.

Procedure II.

Part 1.

If $k > 0$, we have some crossings on \overline{ab} . Now, since $\overline{a_0 a_1}$ is the overstrand (or understrand as the case may be) of a crossing, we cut $\overline{a_0 a_1}$ at some point between a_0 and a_1 , which we shall call P . We then pull the two loose ends apart and over (or under if $\overline{a_0 a_1}$ is the understrand) all other arcs so that the endpoint coming from a_1 lies on A and the other lies on B (with the remainder of both strands still lying inside S). If there are any other arcs that have been pulled up to A (i.e., from procedure I), we place the end point in question to the left of all such points on A , and likewise for the end point pulled to B . Label the point pulled to A by P' and the point pulled to B by P'' . A simple matter of ‘tightening’




Figure 50: *Cutting the link at P , and then rejoining for Procedure I.*

ensures that $\overline{a_0 P''}$ and $\overline{a_1 P'}$ have no local minima or maxima. After this, a_0 is no longer a local minimum, and a_1 is now the minimum of the increasing arc $\overline{a_1 b}$. We then join P'' to P' by an oriented arc lying outside S , that begins at P'' , travels around and to the right of S and any other arcs that we have added outside S , then joins P' from above. Figure 51 illustrates this.

Part 2.

We now repeat part 1 for the increasing arc $\overline{a_1 b}$, which has k points marked on it. Repeating part 1 for $\overline{a_2 b}$, $\overline{a_3 b}$ etc we are eventually left with an increasing arc $\overline{a_{k-1} b}$ with only two points marked on it. One further application of part 1 eliminates b as a local maximum and a_{k-1} as a local minimum, without introducing any new local minima or maxima inside S .

(end of procedure II)

Clearly, after applying procedure I or II, the resulting collection of arcs is equivalent (as a link) to L . This is because we can contract the elongated arc $\overline{a_0 P'' P' a_1}$ back to the little arc $\overline{a_0 a_1}$ that it originally was, since all we did was stretch it over (or under) the rest of L . This contraction is an ambient isotopy, so the two are equivalent.

Recall that, essential to procedures I and II is the fact that if there are any local minima in our diagram, then we can find an increasing arc. However, we will soon be required to apply procedure I and/or II again, to our now modified




Figure 51: *Cutting the link at P , and then rejoining for Procedure II part 1.*

diagram. So we must ensure that our method of finding an increasing arc is still valid after already having applied procedure I or II. That is, if there is another local minimum a remaining in S , then by moving along L from a in the direction of orientation we must come to a local maximum in S . So assume we do not. That is, suppose our method fails, whereby there is a local minimum a remaining in S , but by moving along L from a in the direction of orientation, we do not arrive at any local maximum in S . But since a is a local minimum, then as we travel along our curve away from a the curve must be increasing. Since our curve is compact, we must eventually reach a local maximum. We know this maximum is not in S , so before we reach it we must have left S . But the portion of the curve we are following is increasing, and thus must have passed through the top edge A of S (we know we couldn't have passed through a side edge because procedures I and II only move arcs to pass through the top or bottom edges of S). However, after applying procedures I or II, all arcs in S that meet A have orientation *away* from A and towards B . That is, for any endpoint placed on A , the remainder of the arc is oriented away from A back into S . But the curve we are following meets A , and is oriented towards A . This is a contradiction, and so cannot possibly occur. Thus our method of finding an increasing arc does not fail.

So, given an increasing arc \overline{ab} , by applying one of procedure I or II (whichever is appropriate) we reduce each of the number of minima and maxima inside S by 1. But our original diagram D had only a finite number of minima/maxima

in it (say m of them). So by applying a sequence of procedures I/II of length at most m , we obtain a new diagram of our link L , say D' , such that there are no minima or maxima inside S . We will denote the subset of D' lying in S by C .

We now prove that what is inside S is the diagram of a braid. We know that S contains part of the link diagram D' , with no local minima or maxima. Also, each arc in C begins at A and ends at B , and is oriented from A to B (there are no closed arcs in C , for if there were then C would have a local minimum and a local maximum, which is impossible). The arcs in C have a finite number of crossing points, and since there are finitely many of these arcs, we have that C is in fact the diagram of some n -braid β (for some $n \in \mathbb{N}$). Hence, by the way that we have constructed D' (i.e., with the arcs closing up around S), we have that D' is a diagram for the closure of the braid β . So the closure $\tilde{\beta}$ of β is equivalent (as a link) to L , and the proof is complete. \square

In Figure 52 we perform the entire algorithm for finding L as the closure of some braid β , where L is our example link from Figure 48. Notice the small lump in our diagram for L . We could smooth this lump out, and hence make the algorithm shorter for this example. However, we choose not to, in order to illustrate the fact that the algorithm is valid for *any* diagram for the link L , provided it satisfies the conditions stated at the beginning of the proof of Theorem 6.5 .

6.3 Markov's theorem

Given $\beta \in \mathbf{B}_n$ and $\beta' \in \mathbf{B}_m$, we now wish to investigate the conditions on β and β' required to give $\tilde{\beta} \approx_1 \tilde{\beta}'$. Notice how we can speak of equality between $\tilde{\beta}$ and $\tilde{\beta}'$ since they are both links and thus comparable (even though $\beta \in \mathbf{B}_n$ and $\beta' \in \mathbf{B}_m$, with n and m not necessarily equal). We will ultimately state Markov's theorem, which gives necessary and sufficient conditions for two braids to have the same closure.

We begin by defining the Markov moves on a braid β . Eventually we will show by some simple geometric arguments that these moves preserve the closure of β .

Definition 6.6.

The **Type 1 Markov move** on an n -braid β , denoted by M_1 , replaces β by a conjugate $\gamma\beta\gamma^{-1}$, where γ is any n -braid. The inverse move to M_1 is denoted by M_1^{-1} , and replaces β by a conjugate $\gamma^{-1}\beta\gamma$, where γ is any n -braid.

Essentially, M_1 and M_1^{-1} are the same move, as the inverse to conjugation is also conjugation.




Figure 52: *The entire algorithm for writing an oriented link as the closure of a braid, done for our example link from Figure 48.*




Figure 53: The 4-braid $\beta = \sigma_2^{-1}\sigma_1^{-1}\sigma_3^{-1}\sigma_2^{-1}\sigma_3\sigma_1^{-1}\sigma_3\sigma_2\sigma_1\sigma_3^{-1}$ whose closure gives the oriented link L . We have re-drawn β to make the crossings easier to recognise.

Definition 6.7.

The **Type 2 Markov move** on an n -braid β , denoted by M_2 , replaces β by the $(n + 1)$ -braid $\beta\sigma_n$ or $\beta\sigma_n^{-1}$, where here we view β as a word over the Artin generators $\{\sigma_1^{\pm 1}, \dots, \sigma_{n-1}^{\pm 1}\}$ and thus $\beta\sigma_n^{\pm 1}$ makes sense as an $(n + 1)$ -braid. The inverse move to M_2 is denoted by M_2^{-1} , and replaces an $(n + 1)$ -braid $\beta\sigma_n^{\pm 1}$ by the n -braid β , where β is a word on the Artin generators $\{\sigma_1^{\pm 1}, \dots, \sigma_{n-1}^{\pm 1}\}$ only (i.e., $\sigma_n^{\pm 1}$ does not occur in the word β).

Definition 6.8.

Two braids $\beta \in \mathbf{B}_n$ and $\beta' \in \mathbf{B}_m$ are said to be **Markov equivalent**, denoted $\beta \sim_M \beta'$, if there exists a finite sequence of Markov moves that transform β into β' .

It is simple to see that Markov equivalence is an equivalence relation. What is more interesting is that it is a very strong form of equivalence, and is all we need to look at to determine if the closures of two braids are the same.

Theorem 6.9. (*Markov's theorem*)

Let $\beta \in \mathbf{B}_n$ and $\beta' \in \mathbf{B}_m$. Then $\tilde{\beta} \approx_l \tilde{\beta}'$ (as oriented links) if and only if $\beta \sim_M \beta'$.

Proof.

We will prove one direction of the theorem; that $\beta \sim_M \beta'$ implies $\tilde{\beta} \approx_l \tilde{\beta}'$, and refer the reader to [4], pp 148-163 for a proof of the other direction. We will show that a Markov equivalence preserves the closure of a braid. We begin by fixing some $n \in \mathbb{N}$ and some $\beta \in \mathbf{B}_n$.




Figure 54: A diagram for $g\beta g^{-1}$, where A is a diagram for g and B is a diagram for β .




Figure 55: A diagram for the closure of $g\beta g^{-1}$.

M_1 :

Choose any $g \in \mathbf{B}_n$. Then $M_1(\beta) = g\beta g^{-1}$. Let B be a diagram for β , and A be a diagram for g . Then ABA^{-1} is a diagram for $M_1(\beta)$, as in Figure 54. And thus a diagram for the closure of $M_1(\beta)$ is given by closing the diagram for ABA^{-1} , as shown in Figure 55. However, it is immediately obvious that we can slide the sub-diagram A up, then around to the right, then down, then around to the left, then up again, to meet up with the sub-diagram A^{-1} . Figure 56 illustrates how this is done. But this move is an ambient isotopy, and thus is equivalent to the original closure (as a link). Hence the closure of $M_1(\beta)$ is equivalent to $\tilde{\beta}$ (as a link). Further, setting $h = g^{-1}$ shows that $M_1^{-1}(\beta)$ is equivalent to $\tilde{\beta}$ (as a link), since M_1^{-1} is also conjugation by an element of \mathbf{B}_n .

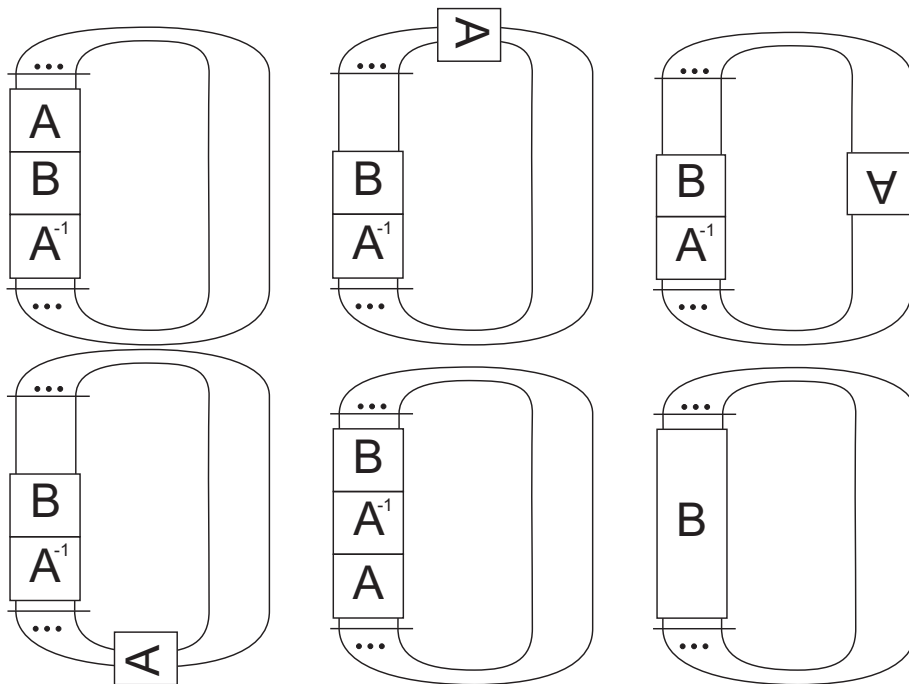


Figure 56: *Sliding the sub-diagram A around the diagram for the closure of $g\beta g^{-1}$. Observe that, once we have slid A all the way around to meet up with A^{-1} , the two sub-diagrams then cancel. So we are left with the closure of β .*




Figure 57: A diagram for the closure of $\beta\sigma_n$.




Figure 58: Untwisting a loop in the link $(\tilde{\beta}\sigma_n)$ to get the link $\tilde{\beta}$.

M_2 :

Given $\beta = \sigma_{i_1}^{\epsilon_1} \dots \sigma_{i_k}^{\epsilon_k} \in \mathbf{B}_n$, we now form the $(n+1)$ -braid $\sigma_{i_1}^{\epsilon_1} \dots \sigma_{i_k}^{\epsilon_k} \sigma_n^\epsilon$. Let B be a diagram for β . Then Figure 57 gives a diagram for the closure of $\beta\sigma_n$ (we have set $\epsilon = 1$, the case for $\epsilon = -1$ is analogous). As can be seen in Figure 58, the innermost loop in the diagram for $(\tilde{\beta}\sigma_n)$ can be untwisted by an ambient isotopy, so we are left with a diagram for $\tilde{\beta}$. Thus these two links must be equivalent, and hence the closure of $M_2^{-1}(\beta\sigma_n^{\pm 1})$ is equivalent (as a link) to $\tilde{\beta}$. Reversing the order of the diagrams in Figure 58 also shows that the closure of $M_2(\beta)$ is equivalent to the closure of β .

Thus the closure of a braid remains unchanged after performing one (and hence finitely many) Markov moves or their inverses. Hence $\beta \sim_M \beta'$ implies $\tilde{\beta} \sim_l \tilde{\beta}'$. □

Combining the results of Alexander and Markov, we can compare two links L and L' as follows: Find braids β, β' whose closures are equivalent to L, L' respectively. Then, check if β and β' are Markov equivalent. From the above theorems, we

conclude that $L \approx_l L'$ if and only if $\beta \sim_M \beta'$. However, there is a slight problem in this method, for at present there is no known algorithm to determine if two braids are Markov equivalent (see [4], p 148). Still, for certain cases, it may be convenient to use this test to check if L and L' are equivalent. But we cannot yet use it to compare two links in general.

7 Conclusion

7.1 Summary

The theory of braids extends far beyond the material covered in this work. However, we have attempted to give an elementary, yet thorough, introduction to braids. We have established clear definitions of a braid and braid equivalence, as well as developed a systematic way of drawing braids. We have shown that the set of equivalence classes of n -braids can be viewed as a group (with an appropriate operation), and have found a finite presentation of this group \mathbf{B}_n . We then moved on to derive some interesting properties of \mathbf{B}_n . We found the centre of \mathbf{B}_n , showed \mathbf{B}_n was an infinite group, and also showed that \mathbf{B}_m could be viewed as a subgroup of \mathbf{B}_n for any $m \leq n$. We went on to define some braid invariants, and found an interesting subgroup \mathbf{P}_n of \mathbf{B}_n that had an easy interpretation as a set of geometric braids. We also found some quotients of \mathbf{B}_n that could be easily understood geometrically, and stated the conditions required for these to be finite.

We have given a full solution to the word problem on \mathbf{B}_n , by describing a unique normal form for each element. We also sketched a solution to the conjugacy problem. We concluded by looking at the correspondence between braids and links, and defined a way to turn a braid into an oriented link. We proved a theorem by Alexander that gave us an algorithm for viewing any oriented link as the closure of a braid, and also a theorem by Markov that gave necessary and sufficient conditions for two braids to have equivalent closures.

7.2 Further remarks

There are many other interesting ideas in braid theory that we were not able to cover. These include not only additional properties of the braid group \mathbf{B}_n , but also some applications, as well as an extension of the concept of braids to other spaces. We briefly mention some of these now.

Additional properties of \mathbf{B}_n :

It has been shown that the braid group \mathbf{B}_n is linear (see [12]), automatic (see [6], pp 195-201), and torsion-free (see [2], p 80).

Applications:

The complexity of solving the conjugacy problem on \mathbf{B}_n is being considered as a possible basis for a public key cryptography system. However, as with most encryption systems, there is currently no proof that the conjugacy problem cannot be solved in less than exponential time. See [2], pp 71-72 for further discussion on this.

Braids in other spaces:

The concept of an n -braid can be extended to different spaces. It is easy to imagine our cube \mathbb{D} as a thickened disk $D \times I$, and our braid strings as paths in $D \times I$ starting on $(D, 0)$ and ending on $(D, 1)$. Now, let \mathbb{F} be any surface (i.e., compact 2-manifold). Then take the ‘thickening’ of \mathbb{F} , $\mathbb{F} \times I$, labeling n ordered points on $(\mathbb{F}, 0)$ and n ordered points on $(\mathbb{F}, 1)$. Then add n non-intersecting paths in $\mathbb{F} \times I$, each starting at some unique labeled point on $(\mathbb{F}, 0)$ and ending at some unique labeled point on $(\mathbb{F}, 1)$. Then this collection of paths (which we shall now call strings) is an n -braid for \mathbb{F} . These have very similar properties to our standard braids. We can also define a group of such braids, denoted $\mathbf{B}_n(\mathbb{F})$, which turns out to be a quotient of B_n . See [4], pp 199-205 for a more thorough introduction to braid groups of general surfaces.

Braids via configuration spaces:

There exists a much more general way to view braids, where we look at the motion of points on the disk D . We denote the n -fold product space $\prod_{i=1}^n D$ by P , and form the subspace $F_n(D) := \{(z_1, \dots, z_n) \in P \mid z_i \neq z_j \text{ if } i \neq j\}$ consisting of points in P with distinct co-ordinates. We can then define an equivalence relation \sim on $F_n(D)$, with $(z_1, \dots, z_n) \sim (z'_1, \dots, z'_n)$ if they differ by a permutation of co-ordinates. We define the pure braid group \mathbf{P}_n of D by the fundamental group $\pi_1(F_n(D))$. Similarly, we define the (full) braid group \mathbf{B}_n of D by $\pi_1(F_n(D)/\sim)$. In essence, what we are doing is fixing n points in D , then looking at how these points behave under perturbations. It can be shown that these definitions of \mathbf{P}_n and \mathbf{B}_n coincide with our previous definitions of them. In addition to this, we can generalise the above idea and replace D by any manifold \mathbb{M} , to obtain the pure braid group and braid group of \mathbb{M} . And when \mathbb{M} is a surface, this coincides with our idea of braid groups on surfaces as outlined above. See [3], pp 3-35 for a more comprehensive introduction to configuration spaces.

References

- [1] E. Artin, *Theorie der Zöpfe*, Hamburg Abh **4** (1925), 47-72.
- [2] J. Birman and T. Brendle, *Braids: A Survey*, arXiv:math.GT/0409205 v2, 2 December 2004.
- [3] J. Birman, *Braids, Links and Mapping Class Groups*, Ann. of Math. Studies 82, Princeton University Press (1974).
- [4] K. Murasugi and B. Kurpita, *A Study of Braids*, Kluwer Academic Publishers, (1999).
- [5] K. Murasugi, *Knot Theory and Its Applications*, Birkhäuser, (1996).
- [6] D. Epstein *et al*, *Word Processing in Groups*, Jones and Bartlett Publishers Inc, (1992).
- [7] F. Garside, *The Braid Group and Other Groups*, Quart. J. Math Oxford **20** (1969), 235-254.
- [8] C.Adams, *The Knot Book: An Elementary Introduction to the Mathematical Theory of Knots*, W. H. Freeman and Company, (1994).
- [9] E. ElRifai and H. Morton, *Algorithms for Positive Braids*, Quart.J. Math. Oxford Ser (2), **45** (180) (1994), 479-497.
- [10] J. Birman, D. Long and J. Moody, *Finite-Dimensional Representations of Artin's Braid Group*, in "The Mathematical Legacy of Wilhelm Magnus", Contemporary Math **169** (1994), 123-132.
- [11] V. Gebhardt, *A New Approach to the Conjugacy Problem in Garside Groups*, preprint arXiv:math.GT/0306199 v2, 21 October 2003.
- [12] D. Krammer, *Braid Groups are Linear*, arXiv:math.GR/0405198 v1, 11 May 2004.