Non-Isotopic Symplectic Surfaces in Products of Riemann Surfaces

by

Christopher Hays

A thesis
presented to the University of Waterloo
in fulfilment of the
thesis requirement for the degree of
Master of Mathematics
in
Pure Mathematics

Waterloo, Ontario, Canada, 2006

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Let Σ_g be a closed Riemann surface of genus g. Generalizing Ivan Smith's construction, for each $g \geq 1$ and $h \geq 0$ we construct an infinite set of infinite families of homotopic but pairwise non-isotopic symplectic surfaces inside the product symplectic manifold $\Sigma_g \times \Sigma_h$. In particular, we achieve all positive genera from these families, providing first examples of infinite families of homotopic but pairwise non-isotopic symplectic surfaces of even genera inside $\Sigma_g \times \Sigma_h$.

Acknowledgements

This thesis is a result of many hours spent working with my supervisor Professor Doug Park. Both this thesis and my understanding of the field in general are much stronger because of this.

Secondly, I am indebted to Professors Ian Goulden, John Lawrence, John D. McCarthy and David McKinnon, as well as Mainak Poddar for the time that they have spent answering questions posed to them.

Thirdly, I wish to acknowledge the many people who put up with my incessant ramblings on the subject, the least of whom is not Martha Yip.

Contents

List of Figures

Chapter 1

Introduction

1.1 Background

Given a symplectic 4-dimensional manifold (M, ω) and a second homology class $a \in H_2(M; \mathbb{Z})$, one can ask for a classification of the isotopy classes of symplectic surfaces that represent a. Problems along these lines are typically labelled as the 'symplectic isotopy problem'. One reason that the symplectic isotopy problem is of interest is that it contrasts between symplectic and Kähler topology.

Given a complex surface S (real dimension 4) and homology class $a \in H_2(S; \mathbb{Z})$, we can consider the complex curves of S that represent a. Let $\alpha \in H^2(S; \mathbb{Z})$ be Poincaré dual to a. Any complex curve C of S representing a will be the zero set of a holomorphic section of a line bundle whose first Chern class is α. It therefore suffices to study the preimage of α under $c_1 : \check{H}^1(S, \mathcal{O}^*) \to H^2(S; \mathbb{Z})$. This set naturally admits the structure of a variety. Moreover, when this preimage is non-empty the subset corresponding to singular surfaces is a proper subvariety and therefore has positive codimension. Thus, there are at most finitely many complex surfaces representing a.

Additionally, the Lefschetz Theorem on $(1, 1)$-classes implies that the above c_1 has image $H^{1,1}(S; \mathbb{C}) \cap H^2(S; \mathbb{Z})$, and so these are the only homology classes that admit complex representatives. This argument is presented in greater depth in Chapter 1 Section 2 of [15].

Note that the generalized Thom conjecture (proved by Morgan, Szabó and Taubes in [27]) states that if a admits a complex representative, then the complex representative has minimal genus among all surfaces representing a. In particular complex surfaces representing the same homology element admit the same genus.

When working in the symplectic setting, the answer to this question is not as clean. Let $a \in H_2(M;\mathbb{Z})$. We again have that symplectic representatives of a have minimal genus among all surfaces representing the homology class (this was proved by Ozsváth and Szabó in [28]). However, there is no known method to determine the number of isotopy classes of symplectic surfaces representing a. It was shown by Siebert and Tian in [32] that there do exist elements of $H_2(\mathbb{CP}^2, \mathbb{Z})$ that are represented by a unique isotopy class of symplectic surfaces. In contrast to this, it was shown by Fintushel and Stern in [12] that certain simply connected 4-manifolds admit infinitely many homology classes that are each represented by an infinite number of symplectic tori. This naturally leads to the question of which genera can be obtained as the genera of symplectic surfaces representing a given second homology class when there are infinitely many isotopy classes of such representatives.

It was then proved by Smith in [33] that for each odd $g \neq 3$ there exist an $h \in \mathbb{N}$ such that $\Sigma_h \times S^2$ admits a homology class that is represented by infinitely many non-isotopic symplectic surfaces of genus g. This result was improved upon by Park, Poddar and Vidussi in [30], who showed that for each $g \geq 1$ there exists a simply connected symplectic 4-manifold with a second homology class that is represented by infinitely many non-isotopic surfaces of genus g.

The purpose of this thesis is to generalize Smith's construction so that we can realize all genera greater than 0 as the genus of symplectic surfaces representing a given homology class for which there exists an infinite number of non-isotopic symplectic representatives. We generalize Smith's construction in two ways. First, we work with more general braid groups $\mathrm{Br}_n(\Sigma_h)$ for all $h \geq 0$. This is not required to obtain the even genera, but rather done so as to generalize the symplectic manifolds and homology classes for which we can construct an infinite number of non-isotopic symplectic representatives; we can then work with the symplectic manifolds $\Sigma_g \times \Sigma_h$ for $g \geq 1$ and $h \geq 0$. Second, in place of studying the hyper-elliptic 2-fold branched cover of S^2, we study all cyclic n-fold branched covers of all Riemann surfaces.

When studying the above symplectic manifolds, we are no longer considering simply-connected manifolds. In this setting we are introduced to a new phenomenon: two homologous surfaces can now fail to be isotopic because they are not even homotopic. We have therefore reformulated the problem to consider the stronger case of homotopic but non-isotopic surfaces. When doing so, we prove the following Theorem.

Theorem 1. *Let g, h, p be integers such that either $g \geq 1$, $h = 0$, $p \geq 4$ or $g \geq 1$, $h \geq 1$, $p \geq 2$. Considering the product symplectic form on $\Sigma_g \times \Sigma_h$, the homology class $p[\Sigma_g \times \Sigma_h] \in H_2(\Sigma_g \times \Sigma_h)$ is represented by infinitely many homotopic but non-isotopic surfaces of genus $p(g-1)+1$.*

Note that in Theorem 70 we will actually consider more general homology classes that will be described later.

1.2 Outline

This thesis is organized as follows. In Chapter 2 we introduce the notion of fibre bundles and state all results used later, with the exception of the theory of symplectic vector bundles.

In Chapter 3 we provide the underlying theory of braid groups and mapping class groups of a manifold. The first important theorem of this chapter is Theorem 37, which defines a specific map from the n^{th} braid group to the n-punctured mapping class group of a Riemann surface. The second important theorem is Theorem 48, which relates the n-punctured mapping class group of an oriented manifold M with the ζ-mapping class group of a manifold N and automorphism ζ such that taking the quotient of N by ζ induces an n-fold cyclic branched cover of N onto M. This theorem is the main tool used to generalize Smith's construction to all cyclic n-fold branched covers.

In Chapter 4 we introduce the notion of a symplectic manifold. There are two major theorems proven in this Chapter. The first theorem is the adjunction formula, which provides an equation that demonstrates a relation between the self-intersection of a symplectic surface with the genus of the surface. We use this later to determine the genus of the constructed surfaces. The second theorem is a special case of a theorem due to Gompf that defines a symplectic summing process of a manifold along two disjoint symplectomorphic submanifolds of codimension 2. We are concerned with the case when the two codimension 2 submanifolds are contained in different components of the ambient manifold, and each has trivial normal bundle.

In Chapter 5 we construct the families of surfaces mentioned in Theorem 1.

Note that the citations appearing beside theorems within this thesis are not meant to provide the original source of the theory, but only the reference that I used while learning the theory.

Chapter 2

Fibre Bundles and Branched Covers

2.1 Introduction

The theory of fibre bundles plays multiple roles within this thesis; they will receive mention in all upcoming chapters. As such, the purpose of this chapter is to coherently organize the theorems that will be used later. This chapter is organized as follows. In Section 2.2 we provide the definition of a fibre bundle and make note of the long exact sequence of homotopy groups that arises from a fibre bundle. The theory described in this section is developed from [22] and [34]. In Section 2.3 we consider the special case of vector bundles, where each fibre has a (real) vector space structure. In Section 2.4 we consider branched covers, which are fibrations with nicely controlled singularities.

2.2 Fibre Bundles

There is an intuitive notion that motivates the definition of a fibre bundle: we wish to consider spaces formed by gluing a copy of some preferred space (the fibre) over each point in some other space (the base space). With this picture in mind, we define a *pre-bundle* over a space M to be a continuous surjective function $\pi : E \to M$ whose fibres $\pi^{-1}(p)$ are pairwise homeomorphic. For ease of notation, when the base space is understood we will also denote a pre-bundle

$\pi : E \to M$ as E_π.

The easiest examples of pre-bundles are direct products of two spaces: given spaces M and F, the canonical projection $M \times F \to M$ defines a pre-bundle with base space M and fibre F. We will refer to these pre-bundles as *trivial pre-bundles.*

The concept of a pre-bundle admits certain pathologies. Consider the set

$$S = \{(x, y, 0) \in \mathbb{R}^3 \mid x > 0\} \cup \{(x, 0, z) \in \mathbb{R}^3 \mid x \leq 0\}$$

with the subspace topology. There exists a map $S \to \mathbb{R}$ defined by projection onto the first coordinate and the pre-image of any point in $\mathbb{R}$ under this map is homeomorphic to $\mathbb{R}$. However, the fibre over $(0, 0, 0)$ is not attached to the fibres over small positive numbers in a desirable manner. To remedy this pathology, we demand that every point in the base space admits a neighbourhood about which the pre-bundle is 'nice'. Lastly, we demand that the 'nice' neighbourhoods are compatible. These stipulations provide the definition of a fibre bundle.

Definition 2. *A **fibre bundle** over a manifold M with fibre F is a pre-bundle $\pi : E \to M$ with fibres homeomorphic to F satisfying the following two conditions.*

- *There exists an open cover $\{U_\alpha\}_{\alpha \in \Lambda}$ of M and a family of homeomorphisms $\{h_\alpha : \pi^{-1}(U_\alpha) \to U_\alpha \times F\}_{\alpha \in \Lambda}$ such that the following diagram commutes:*

$$\begin{array}{ccc} \pi^{-1}(U_\alpha) & \xrightarrow{h_\alpha} & U_\alpha \times F \\ & \searrow^{\pi} & \downarrow \\ & & U_\alpha \end{array}$$

  where the vertical map is the canonical projection.

- *For each pair $\alpha, \beta \in \Lambda$ the map $g_{\beta\alpha} : U_\alpha \cap U_\beta \to Aut(F)$ defined by $g_{\beta\alpha}(x)(f) = h_\beta \circ h_\alpha^{-1}(x, f)$ is continuous.*

Note that the commutativity of the diagram in the first axiom is equivalent to demanding that h_α maps the fibre over $x \in U_\alpha$ in $\pi^{-1}(U_\alpha)$ homeomorphically onto the fibre over x in $U_\alpha \times F$. In the second axiom $\mathrm{Aut}(F)$ is the set of homeomorphisms of F equipped with the compact-open topology. We refer to the collection $\{h_\alpha : \pi^{-1}(U_\alpha) \to U_\alpha \times F\}_{\alpha \in \Lambda}$ as a *local trivialization* of E_π. We refer to the maps $g_{\beta\alpha}$ as *coordinate transformations* or *cocycles* of the bundle. We will refer to the codomain of the cocycles as the *structure group* of the trivialization.

The trivial pre-bundles are fibre bundles that admits a local trivialization consisting of just the identity map of E_π. On the other hand, the second example of a pre-bundle clearly fails to be a fibre bundle since there is no such neighbourhood about 0.

A compatible collection of cocycles is sufficient to fully describe the fibre bundle. Given spaces M and F, consider an open cover $\{U_\alpha\}_{\alpha\in\Lambda}$ of M and a collection of continuous maps $\{g_{\beta\alpha} : U_\alpha \cap U_\beta \to G\}_{(\alpha,\beta)\in\Lambda\times\Lambda}$ where G is a subgroup of the automorphism group of F. Suppose furthermore that these g's satisfy

$$g_{\gamma\beta}(x)g_{\beta\alpha}(x) = g_{\gamma\alpha}(x) \tag{2.1}$$

for all $x \in U_\alpha \cap U_\beta \cap U_\gamma$. The g's then necessarily satisfy the identities

$$g_{\alpha\alpha} = id \quad \text{and} \quad g_{\beta\alpha} = g_{\alpha\beta}^{-1} \tag{2.2}$$

for all $\alpha, \beta \in \Lambda$. Regarding Λ as being equipped with the discrete topology, consider the subset $T \subset M \times F \times \Lambda$ consisting of all points (x, f, α) such that $x \in U_\alpha$. Alternatively we can characterize T as the union

$$\bigsqcup_{\alpha\in\Lambda} U_\alpha \times F \times \{\alpha\}.$$

equip T with the binary relation

$$(x_1, f_1, \alpha) \sim (x_2, f_2, \beta) \quad \text{if} \quad x_1 = x_2 \quad \text{and} \quad g_{\beta\alpha}(x_1)(f_1) = f_2.$$

It is immediate from (2.1) and (2.2) that this relation is an equivalence relation. Let E denote the space of equivalence classes of $\sim$ with the quotient topology. Define $\pi : E \to M$ by sending $\pi(x, f, \alpha)$ to x. Since π is well-defined and the lift $\tilde{\pi} : \sqcup U_\alpha \times F \times \{\alpha\} \to M$ is continuous, we also have that π is continuous. Moreover, since $\pi^{-1}(x) \cong F$ for all $x \in M$ and each $x \in M$ is contained in one of the U_α's (which provides a local trivialization), E_π is a fibre bundle.

From the above construction we see that the only role that the fibre plays in the construction of a fibre bundle is in understanding of how the transformation group acts on the fibre. In particular, it is possible to replace the fibre with some other space that G acts upon. We will say that two bundles are *associated* to one another if each can be constructed from the other in this manner. One particular example of this is the *associated principal bundle*, where one replaces the fibre with the space G acting on itself by left multiplication.

Note that the structure group G depends upon the choice of open cover of M and cocycles g. In particular, Theorem 6.4.1 of [22] states that given a closed subgroup H of G, one can reduce the structure group to H if and only if for each $\alpha \in \Lambda$ there exists a map $r_\alpha : U_\alpha \to G$ such that the functions $r_\beta^{-1} g_{\beta\alpha} r_\alpha$ have values in H along $U_\alpha \cap U_\beta$.

It is convenient to assume that we have a fixed local trivialization of the bundle, and to implicitly work with local coordinates (x, y) of E where x is a coordinate for M and y is a coordinate for F.

Having defined fibre bundles, we now proceed to define the natural notion of a morphism between two bundles.

Definition 3. *Let E_π and E_ρ be two fibre bundles over M. A morphism between these two bundles is a continuous map $f : E_\pi \to E_\rho$ such that*

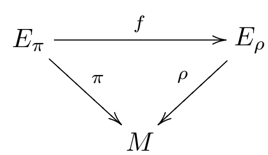

commutes.

Note that the commutativity requirement is equivalent to requiring that the fibre over $x \in M$ in E_π maps to the fibre over x in E_ρ.

The collection of all bundles over a space M together with morphisms between these bundles defines a category $\boldsymbol{Bun}(M)$.

It will occasionally be convenient to loosen the requirement that we have naturally identified the base spaces of two manifolds. We can therefore more generally define a morphism as follows.

Definition 4. *Let $\pi : D \to M$ and $\rho : E \to N$ be two fibre bundles. A morphism between D_π and E_ρ is a pair of continuous maps $f : D \to E$ and $g : M \to N$ such that the following diagram commutes.*

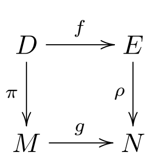

The commutativity requirement here is that the fibre over $x \in M$ maps to the fibre over $g(x)$. Using this broader notion, we can consider the collection of all bundles and all such morphisms to define a category $\boldsymbol{Bun}$.

Rather than working over all topological spaces, we may instead consider only fibre bundles defined within other geometric categories. Indeed, we will primarily consider our bundles to be within the smooth category. When considering smooth fibre bundles we demand that the projection is a submersion and that there exists a local trivialization such that the induced maps $g_{\beta\alpha}$ have their image contained within the subspace of diffeomorphisms of the fibre.

Definition 5. *A* ***section*** *of a bundle $\pi : E \to M$ is a morphism $\sigma : M \to E$ such that $\pi \circ \sigma = id_M$.*

In the case of a trivial bundle $M \times F$, the sections are precisely graphs of continuous functions $M \to F$. Since each fibre intersects the image of a section at a single point, we can identify sections with their images. Thus, a section can be reinterpreted as a submanifold $\sigma(M)$ of E such that each $\sigma(x) \in \sigma(M)$ is contained in $\pi^{-1}(x)$.

The notion of a section naturally generalizes to the notion of a multi-section.

Definition 6. *An* ***n-section*** *of a bundle $\pi : E \to M$ is a submanifold X of E such that the cardinality of the intersection of X with any fibre $\pi^{-1}(x)$ is n.*

Note that an n-section, together with the restriction of π to the n-section, is itself a fibre bundle over M whose fibre is a discrete set of size n. Indeed, the local trivialization of E_π restricts to a local trivialization of the n-section, and since the homeomorphism group of the fibre is discrete the $g_{\beta\alpha}$'s are necessarily continuous.

Fibre bundles with a discrete finite fibre will come up later. These bundles are typically referred to as *covering spaces* of the base space.

It is a classical result that the connected finite covers E_π over a connected manifold, and more generally the fibre bundles whose fibres are merely discrete, are classified by their image $\pi_*(\pi_1(E)) \subset \pi_1(M)$.

Theorem 7. [17] *Let M be path-connected and locally path-connected. There exists a one-to-one correspondence between connected discrete covers E_π of M and subgroups of $\pi_1(M)$ given by pairing E_π with $\pi_*(\pi_1(E)) \subset \pi_1(M)$. More generally, if we only require that π is an $n : 1$ local homeomorphism, there is a* *bijection between such covers and conjugacy classes of subgroups of* $\pi_1(M)$*. In any case, the cardinality of the fibre is equal to the index* $[\pi_1(M) : \pi_*(\pi_1(E))]$.

A proof of this theorem can be found, for instance, on page 63 of [17].

Applications of the of fibre bundles in the next chapter will make use of the following theorem, which demonstrates the relationship between the homotopy groups of the spaces involved within a fibre bundle.

Theorem 8. [34] *Let* $\pi : E \to M$ *be a fibre bundle, with fibre* F*. Fix* $x \in M$ *and* $y \in \pi^{-1}(x)$*. Let* $i : F \hookrightarrow E$ *be the inclusion of* F *as the fibre over* x*. Then, there exists a long exact sequence:*

$$\begin{aligned} \longrightarrow \pi_n(F) \xrightarrow{i_*} \pi_n(E) \xrightarrow{\pi_*} \pi_n(B) \longrightarrow \cdots \\ \cdots \xrightarrow{\pi_*} \pi_1(B) \longrightarrow \pi_0(F) \xrightarrow{i_*} \pi_0(E) \xrightarrow{\pi_*} \pi_0(B). \end{aligned}$$

A proof of this theorem can be found on page 91 of [34].

Note that the last three morphisms are not group morphisms (unless this is a bundle within the category of topological groups, in which case $\pi_0(F)$, $\pi_0(E)$ and $\pi_0(B)$ are actually groups) but rather morphisms of pointed sets. We are therefore considering exactness at these points to be exactness within this category of pointed sets. Note that the image of a group morphism under the forgetful functor to the category of pointed sets has kernel and image equal to the kernel and image of the group morphism.

2.3 Vector Bundles

Vector bundles are fibre bundles where the fibres have a vector space structure that is locally compatible.

Definition 9. *An* ***n-dimensional*** (***real***) ***vector bundle*** *is a fibre bundle* $\pi : E \to B$ *with fibre* $\mathbb{R}^n$ (*considered to be equipped with the natural vector space structure on* $\mathbb{R}^n$) *that admits a local trivialization*

$$\{h_\alpha : \pi^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^n\}_{\alpha \in \Lambda}$$

such that for $p \in U_\alpha$ *the map* $h_\alpha|_{\pi^{-1}(p)} : \pi^{-1}(p) \to \mathbb{R}^n$ *is a vector space isomorphism.*

Given a local trivialization of a vector bundle, the requirement that the h_α's restrict to vector space isomorphisms on each fibre guarantees that the cocycles $g_{\beta\alpha}$ have their image contained within the space of $\mathrm{Aut}(\mathbb{R}^n)$ consisting of linear isomorphisms. Indeed, we can redefine an n-dimensional (real) vector bundle as a fibre bundle whose fibre is $\mathbb{R}^n$ and whose structure group is a subgroup of $GL_n(\mathbb{R})$ with the natural action.

Definition 10. *A fibre bundle morphism $f : E_\pi \to E_\rho$ between two vector bundles is a* ***vector bundle morphism*** *if the restriction of f to each fibre is a linear morphism.*

Let $\boldsymbol{Vect}(M)$ denote the category of vector bundles over M. This category admits much more structure than $\boldsymbol{Bun}(M)$. For example, $\boldsymbol{Vect}(M)$ admits a zero-object: the identity map $id : M \to M$ is a vector bundle whose fibre is the zero-dimensional vector space, and for any other bundle E_π over M there exist unique morphisms $E_\pi \to M_{id}$ and $M_{id} \to E_\pi$. As a consequence, given two vector bundles E_π and E_ρ there exists the zero-morphism $0_{\pi\rho} : E_\pi \to M_{id} \to E_\rho$. This allows for various constructions within $\boldsymbol{Vect}(M)$.

Given local trivializations $\{g_\alpha : \pi^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^n\}_{\alpha\in\Lambda}$ of E_π and $\{h_\beta : \rho^{-1}(V_\beta) \to V_\beta \times \mathbb{R}^m\}_{\beta\in\Gamma}$ of E_ρ, we can locally consider $h_\beta \circ f \circ g_\alpha^{-1} : (U_\alpha \cap V_\beta) \times \mathbb{R}^n \to (U_\alpha \cap V_\beta) \times \mathbb{R}^m$. Since the fibres of E_π and E_ρ over $U_\alpha \cap V_\beta$ are trivially isomorphic to $\mathbb{R}^n$ and $\mathbb{R}^m$ respectively, and since f is continuous, we have a well-defined continuous map $f : U_\alpha \cap V_\beta \to \mathrm{Hom}(\mathbb{R}^n, \mathbb{R}^m)$. Since this map is continuous, f has locally constant rank. Thus, assuming M is connected, the restrictions of f to any two fibres have the same rank.

This fact allows us to associate a number of vector bundles to a vector bundle morphism. Define $K = \{(x, y) \in E_\pi \mid f(x, y) = (x, 0)\}$, i.e. let K equal the union of the kernels of the restrictions of f to each fibre. The restriction of π to K provides a surjection onto M, where each fibre is homeomorphic to $\mathbb{R}^k$ for some fixed k. Moreover, given a local trivialization $\{g_\alpha : \pi^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^n\}_{\alpha\in\Lambda}$ of E_π, the restriction $\{g_\alpha : (\pi|_K)^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^k\}_{\alpha\in\Lambda}$ is a local trivialization of K. We therefore have that K is vector bundle over M. Note that K is naturally the kernel of f: given any vector bundle morphism $h : E_\zeta \to E_\pi$ such that $f \circ h = 0_{\zeta\pi}$, the image of h is necessarily contained within K and we can therefore factor h through K.

The image of a vector bundle morphism, $\mathrm{Im} f$, is also a vector bundle over M. Using the above notation, we have that f has constant rank and any local trivialization of E_ρ restricts to a local trivialization of $\mathrm{Im} f$.

Note that these bundles occur as sub-bundles of either the domain or codomain (sub-bundle meaning that they are subsets such that any local trivialization of the total set restricts to a local trivialization of the subset). We will similarly construct cokernels of a vector bundle morphism later.

Other methods of constructing new vector bundles come from globally performing vector space operations along every pair of fibres of two bundles. As an example of this, consider bundles E_π and E_ρ with fixed local trivializations $\{g_\alpha : \pi^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^n\}_{\alpha\in\Lambda}$ and $\{h_\beta : \rho^{-1}(V_\beta) \to V_\beta \times \mathbb{R}^m\}_{\beta\in\Gamma}$ respectively. We can then define a vector bundle

$$E_\pi \oplus E_\rho = \{(p, q) \in E_\pi \times E_\rho \mid \pi(p) = \rho(q)\}$$

with $(\pi \oplus \rho)(p, q) = \pi(p)$. Note that the preimage of any point $x \in M$ is $\pi^{-1}(x) \oplus \rho^{-1}(x)$.

This bundle admits a local trivialization by restricting the local trivialization $g_\alpha \times h_\beta : \pi^{-1}(U_\alpha) \times \rho^{-1}(V_\beta) \to (U_\alpha \times V_\beta) \times \mathbb{R}^{n+m}$.

This operation also admits a converse: if E_ρ is a sub-bundle of E_π, then there exists a unique vector bundle E_ξ such that $E_\pi = E_\rho \oplus E_\xi$. We will begin proving this by showing that every injective morphism of vector bundles is realized as the injective morphism of a short exact sequence (exact if at every stage the image equals the kernel). We will then show that every short exact sequence splits, proving both existence and uniqueness of the E_ξ bundle.

Proposition 11. [18] *Let $f : E_\rho \to E_\pi$ be an injective vector bundle morphism. There exists a vector bundle E_ξ, and a vector bundle morphism $g : E_\pi \to E_\xi$ such that*

$$1 \longrightarrow E_\rho \xrightarrow{f} E_\pi \xrightarrow{g} E_\xi \longrightarrow 1$$

is exact.

Before proving this proposition, we require the following definition.

Definition 12. *An* ***orthogonal structure*** *on a vector bundle E_π is a continuously varying map from M to the space of positive-definite bilinear forms on the fibres of E_π.*

There are a number of equivalent ways to define the above notion of continuous. The easiest is just to demand that the evaluation map $E_\pi \oplus E_\pi \to \mathbb{R}$ is continuous.

Alternatively, we can define an orthogonal structure to be a reduction of the transformation group to $O(n)$. Suppose that E_π is equipped with an orthogonal structure. Choose a local cover of contractible sets $\{U_\alpha\}_{\alpha\in\Lambda}$ of M. Note that over each U_α, E_π is trivial. Thus, for each α we can choose a set of n sections $s_i : U_\alpha \to \mathbb{R}^n$ such that at each $x \in U_\alpha$, $\{s_1(x), \ldots, s_n(x)\}$ is a basis of the fibre. We can then globally apply the Gram-Schmidt orthonormalization process to achieve orthonormal sections $t_1, \ldots, t_n$. The functions $h_\alpha^{-1} : U_\alpha \times \mathbb{R}^n \to \pi^{-1}(U_\alpha)$ defined by $h_\alpha^{-1}(x, v_1, \ldots, v_n) = \sum_{i=1}^n v_i t_i(x)$ are inverses of new local trivializations over U_α for which it is routine to check that the structure group is $O(n)$ (using Theorem 7.4.1 of [22]).

On the other hand, suppose that E_π is a vector bundle with structure group $O(n)$, and let $\{h_\alpha : \pi^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^n\}_{\alpha\in\Lambda}$ be a local trivialization of E_π that induces cocycles that map into $O(n)$. Fixing a choice of $O(n)$ in $GL_n(\mathbb{R})$ determines a preferred bilinear form on $\mathbb{R}^n$ (the form that makes the columns of $O(n)$ orthonormal). Moreover, if we equip the fibres over U_α with this metric according to the trivialization h_α, the change of trivialization using the cocycle $h_{\beta\alpha}$ does not alter the metric. We therefore have a well-defined orthogonal structure on E_π.

We will now proceed with the proof of Proposition 11.

Proof. We will begin by proving that every vector bundle admits an orthogonal structure. Let $\{h_\alpha : \pi^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^n\}$ be a local trivialization of E_π whose open cover of M is locally finite. We can locally define an orthogonal structure on $U_\alpha \times \mathbb{R}^n$ by assigning the identity matrix to each $x \in U_\alpha$. Consider this as a map $\mu_\alpha : U_\alpha \to GL_n(\mathbb{R})$. Let $\{p_\alpha\}_{\alpha\in\Lambda}$ be a partition of unity that is subordinate to $\{U_\alpha\}_{\alpha\in\Lambda}$. We can then define an orthogonal structure on E_π as $\sum_{\alpha\in\Lambda} p_\alpha \cdot \mu_\alpha$.

Alternatively, we could have proven the existence of an orthogonal structure by proving that we could reduce the structure group to $O(n)$ (using Theorem 7.4.1 of [22]).

Having fixed an orthogonal structure of E_π, we can define a surjective vector bundle morphism $E_\pi \to E_\rho$, which fibrewise is the orthogonal projection of $\pi^{-1}(x)$ onto $\rho^{-1}(x)$. Let $E_\rho^\perp$ denote the kernel of this morphism. Define a bundle morphism $g : E_\pi \to E_\rho^\perp$ to be the orthogonal projection onto $E_\rho^\perp$. Fibrewise, $\operatorname{Ker} g = \operatorname{Im} f$, and so $\operatorname{Ker} g = \operatorname{Im} f$ as vector bundles. This, together with the fact that g is surjective, shows that

$$1 \longrightarrow E_\rho \xrightarrow{f} E_\pi \xrightarrow{g} E_\rho^\perp \longrightarrow 1$$

is exact, which completes the proof. ☐

Using the infrastructure of the previous proof, we can also prove the following.

Proposition 13. [18] *Let*

$$1 \longrightarrow E_\rho \xrightarrow{f} E_\pi \xrightarrow{g} E_\xi \longrightarrow 1$$

be a short exact sequence of vector bundles. Then there exists a vector bundle morphism $s : E_\xi \to E_\pi$ such that $g \circ s = id_{E_\xi}$. Equivalently, $E_\pi \cong E_\rho \oplus E_\xi$

Proof. Fix an orthogonal sequence on E_π. We then have the short exact sequence

$$1 \longrightarrow E_\rho \xrightarrow{f} E_\pi \xrightarrow{g} E_\rho^\perp \longrightarrow 1$$

as above, where $E_\rho^\perp$ is a subbundle of E_π. We therefore have the following commutative diagram.

$$\begin{array}{ccccccccc}
1 & \longrightarrow & E_\rho & \xrightarrow{f} & E_\pi & \underset{\hookleftarrow}{\xrightarrow{g}} & E_\rho^\perp & \longrightarrow & 1 \\
 & & \downarrow{\scriptstyle id} & & \downarrow{\scriptstyle id} & & & & \\
1 & \longrightarrow & E_\rho & \xrightarrow{f} & E_\pi & \xrightarrow{g} & E_\xi & \longrightarrow & 1
\end{array}$$

We therefore have a vector bundle morphism $E_\rho^\perp \to E_\xi$. Using the 5-Lemma, this morphism is an isomorphism. In particular, there exists a vector bundle morphism $E_\xi \to E_\pi$ and $E_\pi = E_\rho \oplus E_\xi$. ☐

As a corollary to the above proposition, we have that the cosummand of E_ρ in E_π is unique up to isomorphism.

Corollary 14. *Let $E_\rho \hookrightarrow E_\pi$ be an injective vector bundle morphism. There exists a vector bundle E_ξ, unique up to isomorphism, such that $E_\pi = E_\rho \oplus E_\xi$.*

An example of such a bundle that will be relevant later is the normal bundle of a submanifold M. Given a manifold M with submanifold X, the restriction of the tangent bundle $\pi : TM \to M$, $\pi|_{\pi^{-1}(X)} : TM|_X = \pi^{-1}(X) \to X$, is a vector bundle over X. Moreover, the tangent space of X at a point p can be naturally embedded within the tangent space of M at p. This determines TX as a subbundle of $TM|_X$. We define the cosummand of TX to be the *normal bundle* NX. Note that in this case, the choice of orthogonal structure in Proposition 11 is simply a choice of Riemannian metric on M.

We will return to studying characterizations of the normal bundle of a submanifold in Section 4.4.

2.4 Branched Covers

Definition 15. *A **d-fold branched cover** is a smooth proper map $\pi : E \to M$ with the following two properties.*

- *There exists a codimension 2 submanifold $X \subset M$ such that the restriction $\pi|_{E\setminus\pi^{-1}(X)} : E \setminus \pi^{-1}(X) \to M \setminus X$ is a covering space with fibre of size d.*
- *Near each $p \in \pi^{-1}(X)$, π is locally modeled by $f : \mathbb{C} \times \mathbb{R}^{n-2} \to \mathbb{C} \times \mathbb{R}^{n-2}$ where $f(z, x) = (z^m, x)$ for some $m \in \mathbb{N}$.*

Note that this definition is occurring within the smooth category, and not the complex category, which is more typical in literature.

We will refer to X as the *branch locus* of π, and the m in the second property as the *branch index* of p. Since $\pi|_{\pi^{-1}(M\setminus X)}$ is a covering space of $M \setminus X$, we can fix a decomposition of $\pi^{-1}(M \setminus X)$ into d sheets. For $p \in \pi^{-1}(X)$ the branch index of p is the number of above sheets whose closures contain p. We can therefore associate a partition λ_x of d to each $x \in X$. We will refer to this X-indexed set of partitions as the *branch data* of the branched cover. When M is a closed oriented 2-manifold, the branch locus is a discrete, and hence finite, subset of M. Let us denote the branch locus as $\{x_1, \ldots, x_k\}$. We can then find a collection of k simple closed curves γ_i that separate x_i from the rest of X. We will further assume that each γ_i is oriented as the boundary component of a disk containing x_i. If we fix a base point p_1 of γ_i, the automorphism of $\pi^{-1}(p_i)$ induced by γ_i is a permutation in S_d that is contained within the conjugacy class indexed by λ_{x_i} (using the disjoint cycle decomposition of a permutation). This can be seen by noting that a point $p \in \pi^{-1}(x_i)$ will induce a cycle in the monodromy permutation of size equal to the branch index of p.

Globally, given a base point $y \in M \setminus X$, there exists a natural group morphism $\rho : \pi_1(M \setminus X, y) \to \mathrm{Aut}(\pi^{-1}(y)) \cong S_d$ defined by sending a loop γ to the permutation induced on the sheets by the endpoints of all lifts of γ to $E \setminus \pi^{-1}(X)$. As we vary y, ρ is only determined up to conjugacy. We will refer to this conjugacy class as the *monodromy representation* of π.

Our first objective in this section will be to study the existence of *cyclic* branched covers over a Riemann surface Σ_g.

Definition 16. *A branched cover $\pi : E \to M$ is* ***cyclic*** *provided that the image of the monodromy representation is cyclic.*

Note that any branched cover over a closed oriented surface M equips E with a preferred orientation. Since removing a finite set of points does not affect the orientation of a manifold, $M \setminus X$ is oriented. We can then pull back the preferred orientation form to an orientation form of $E \setminus \pi^{-1}(X)$, and since we can compactify $E \setminus \pi^{-1}(X)$ to E by adding a finite set of points, we have an orientation form on E. In particular, we have the following proposition.

Proposition 17. *Let $\pi : E \to M$ be a branched covering over a closed orientable surface M. Then E is itself a closed orientable surface.*

We are now in a position to prove necessary criteria for the existence of various branched covers over S^2.

Lemma 18. [24] *Let $X = \{x_1, \ldots, x_k\}$ be a subset of S^2. Choose a set map $p : X \to S_d$ such that $\prod_{i=1}^{k} p(x_i) = 1$ and $\langle p(x_1), \ldots, p(x_k)\rangle$ acts transitively on $\{1, \ldots, d\}$. Then there exists a d-sheeted branched cover over S^2 with branch locus X and monodromy group $G = \langle p(x_1), \ldots, p(x_k)\rangle$.*

Proof. Let $G = \langle p(x_1), \ldots, p(x_k)\rangle \leq S_d$. Fix $x_0 \in S^2 \setminus X$. Fix a preferred set of loops γ_i, $1 \leq i \leq k$, that are equipped with the positive orientation and partition x_i from x_0 and the rest of X. We can homotope the γ_i's to loops with endpoint x_0 such that if one were to follow a small circle about x_0 in the positive direction, the γ_i's would be cyclically ordered in their numeric order. We then have that

$$\pi_1(M \setminus X, x_0) = \langle \gamma_1, \ldots \gamma_k \mid \gamma_1 \cdots \gamma_k = 1\rangle.$$

Define the set map $\{\gamma_1, \ldots, \gamma_k\} \to G$ by sending γ_i to $p(x_i)$. Since $\gamma_1 \cdots \gamma_k = 1$, we can uniquely extend this map to a group morphism $\varphi : \pi_1(M \setminus X, x_0) \to G$.

Fix $l \in \{1, \dots, d\}$. Let M_l be the subgroup of G consisting of permutations that fix l. Then, using Theorem 7, $\varphi^{-1}(M_l)$ determines a d-fold covering $\pi : E' \to S^2 \setminus X$.

Since $\langle p(x_1), \dots, p(x_k)\rangle$ acts transitively on $\{1, \dots, d\}$, E' is connected. We can now extend π uniquely to a continuous map with image S^2. The preimage of γ_i under π defines a finite set of circles (of size equal to the number of cycles in the disjoint cycle decomposition of $p(x_i)$). As we shrink γ_i around x_i, the above circles individually shrink around a puncture within E'. Adding points to these punctures, and mapping them to x_i will then define a branched covering over S^2 with the desired properties. $\square$

In particular, if we associate the d-cycle $(1, \dots, d)$ to each $x_i \in X$, we have the following corollary.

Corollary 19. *Let $X = \{x_1, \dots, x_k\}$ be a subset of S^2. Choose $d \in \mathbb{N}$ such that d divides k. There exists a cyclic d-sheeted branched cover over S^2 with branch locus X.*

An example of such a cover can be constructed as follows. Consider the CW-complex of S^2 consisting of two three sided cells. This complex is provided in Figure 2.1. We can lift this CW-complex to a CW-complex of T^2 as shown in Figure 2.2.

The above criterion for the existence of a branched cover with specific branch data can be extended to branched covers over any oriented closed surface. Fix a branched covering $\pi : \Sigma_k \to S^2$ with the desired branch data. Choose discs D_1, $D_2 \subset S^2$ such that $D_1 \cap D_2 = S^1$ and $X \subset D_1$. In $S^2 \setminus X$, the boundary $S^1 = \partial D_2$ is null-homotopic, and thus it will lift to d disjoint circles. More precisely, D_2 lifts to d disjoint copies of D^2 whose boundaries are the lifts of S^1. We can therefore surger π by removing both D_2 and its preimage, and gluing in a copy of $\Sigma_h \setminus D^2$ to each of these circles. We therefore have a branched cover with base space Σ_h. Using Seifert-van Kampen's Theorem we have that

$$\pi_1(\Sigma_h \setminus X) = \langle \gamma_1, \dots, \gamma_k, \alpha_1, \dots, \alpha_h, \beta_1, \dots, \beta_h \mid \gamma_1 \cdots \gamma_k = \prod_{i=1}^{h} [\alpha_i, \beta_i]\rangle.$$

Since each of the α_i and β_i maps to 1 under the monodromy representation, the images of the monodromy representations of the old and the new branched covers can be identified. We therefore have the following theorem.

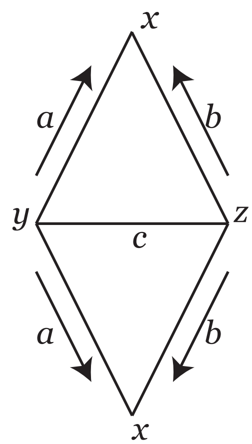

Figure 2.1: A CW-complex of S^2

Theorem 20. *Let $X = \{x_1, \ldots, x_k\}$ be a subset of Σ_h. Choose a set map $p : X \to S_d$ such that $\prod_{i=1}^{k} p(x_i) = 1$. There exists a d-sheeted branched cover over Σ_h with branch locus X and whose monodromy representation has image $G = \langle p(x_1), \ldots, p(x_k)\rangle$. In particular, for all $d \in \mathbb{N}$ such that $d|k$ there exists a cyclic d-sheeted branched cover of Σ_h.*

We now wish to compute the genus of the covering space of a branched cover of Σ_h given the branch data. Given a partition λ, let $\sharp\lambda$ denote the number of summands of λ. Given a branched cover $\pi : \Sigma_k \to \Sigma_h$, we will define the *ramification* of $x \in X$ to be $\nu(x) = d - m$, where d and m are as in Definition 15, and the *ramification index* of π to be $\sum_{x \in X} \nu(x)$.

Given a branched cover $f : E \to M$ with branch locus X, consider a CW-complex of M that extends a CW-complex of X. We then have that the preimage of any cell is a union of subspaces of E that are each homeomorphic to that cell. We can then use these spaces to construct a CW-complex of E. Using this fact, we can determine the genus of the covering space of a branched cover of Σ_h using the branch data.

Proposition 21. *Let $\pi : \Sigma_k \to \Sigma_h$ be a d-fold branched cover with branch locus X of size n. Then $k = d(h-1) + \frac{1}{2}\nu(\pi) + 1$.*

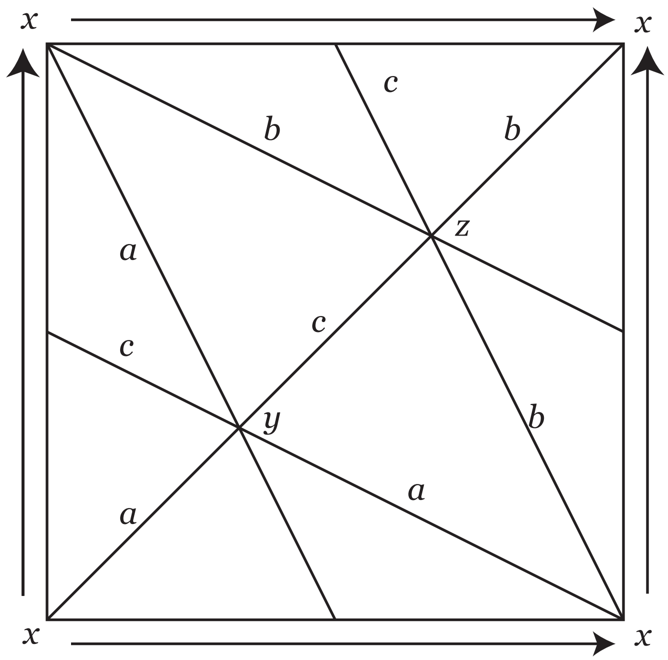

Figure 2.2: A lift of the CW-complex in Figure 2.1

Proof. Choose a CW-complex of Σ_h as above, and lift this to a CW-complex of Σ_k. Each $x \in X$ lifts to $d - \nu(x)$ vertices, and all other cells lift to d cells. We therefore have that

$$\begin{aligned}\chi(\Sigma_k) &= d \cdot (\chi(\Sigma_h) - n) + \sum_{x \in X} (d - \nu(x)) \\ &= d \cdot \chi(\Sigma_h) - \nu(\pi).\end{aligned}$$

In particular,

$$\begin{aligned}k &= \frac{1}{2}(d(2h-2) + \nu(\pi) + 2) \\ &= d(h-1) + \frac{1}{2}\nu(\pi) + 1.\end{aligned}$$ □

When the branched cover is cyclic, the above formula simplifies to the following.

Corollary 22. *Let $\pi : \Sigma_k \to \Sigma_h$ be a cyclic d-fold branched cover with branch locus of size n. Then $k = hd + \frac{1}{2}(2-n)(1-d)$.*

In Section 4.4 we will return to the study of the existence of branched covers, with a focus on branched covers over closed oriented 4-dimensional manifolds.

Chapter 3

Mapping class groups and Braid Groups

3.1 Introduction

The purpose of the first three sections of this chapter is to introduce the braid groups and mapping class groups of manifolds, with particular interest in these groups for the Riemann surfaces.

Section 3.2 motivates the modern definition of the braid groups by first introducing the classical definition. Since the classical braid groups are presented here only to provide motivation and intuition for the modern definition, we refrain from providing any proofs in this setting and from proving the equivalence between the classical and modern definitions for the Euclidean plane. The theorems involving the construction of the classical braid groups are proved in [1] and [2]. A proof showing the equivalence of the two definitions can be found in [5].

Section 3.3 introduces the mapping class groups of a space. Particular attention is paid to morphisms between mapping class groups and braid groups that are obtained by studying particular fibre bundles.

Section 3.4 collects results that provide either a presentation or a set of generators for these groups.

In Section 3.5 we introduce the notion of the ζ-mapping class group, which consists of mapping classes that commute with a fixed mapping class ζ. These groups are generalizations of the hyper-elliptic mapping class group studied by Birman in [3]. The hyper-elliptic mapping class group was studied in order to prove results about the double branched cover obtained from a hyper-elliptic map of a Riemann surface. In this section, we extend these results to all cyclic branched covers.

3.2 The Braid Groups

The braid groups of the Euclidean plane $\mathbb{R}^2$ were introduced by E. Artin in [1] with the intent of characterizing methods for intertwining n strings together. This leads to a very natural characterization of these groups. Let (x, y, t) be coordinates for $\mathbb{R}^2 \times I$, where $x, y \in \mathbb{R}$ and $0 \leq t \leq 1$. Fix points $x_i = (i, 0, 0)$ and $y_i = (i, 0, 1)$ for $i = 1, \ldots, n$. We define a *classical braid representative* to be an embedding $\beta : \sqcup_{i=1}^{n} I \hookrightarrow \mathbb{R}^2 \times I$ such that:

- each interval of $\sqcup I$ has boundary $\{x_i, y_j\}$ for some $1 \leq i, j \leq n$,
- for each $t \in I$, $|\beta^{-1}(\mathbb{R}^2 \times \{t\})| = n$.

A *classical n-braid*, then, is an isotopy class of embeddings $\sqcup_{i=1}^{n} I \hookrightarrow \mathbb{R}^2 \times I$ that contains a such an embedding. Given a braid β, we will denote representatives of β by $(\beta_1, \ldots, \beta_n)$, where β_i is the restriction of the representative to the interval starting at x_i. We will refer to the map β_i as the i^{th} *strand* of β. Given a preferred splitting $\mathbb{R}^2 = \mathbb{R}^1 \times \mathbb{R}^1$, we say that a representative of β is in *general position* provided that under the projection $\mathbb{R}^2 \times I \to \mathbb{R}^1 \times I$, we have at most two strands intersect at a given point, and if two strands do intersect then they do so transversely. Every braid admits a representative that is in general position. This allows us to represent braids as diagrams in $\mathbb{R}^1 \times I$. See Figure 3.2 for an example of such a diagram, and the diagram of an embedding that is not a braid according to the second requirement placed upon braid representatives.

The classical braids admit a binary operation by aligning one braid on top of the other, as in Figure 3.2. This operation admits an identity I_n consisting of n vertical strands. Additionally, one can prove inductively that any braid β admits an inverse, which can be expressed as the 'mirror' of β. An example of this is found in Figure 3.3.

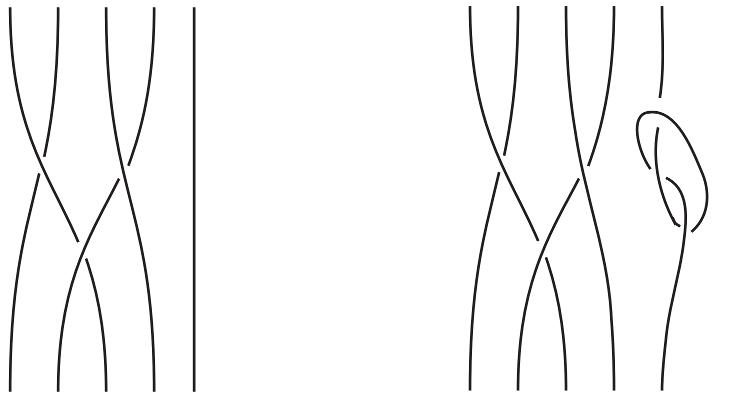

(a) A 5-braid

(b) An embedding without a braid representative

Figure 3.1: Embeddings $\sqcup I \to E^2 \times I$

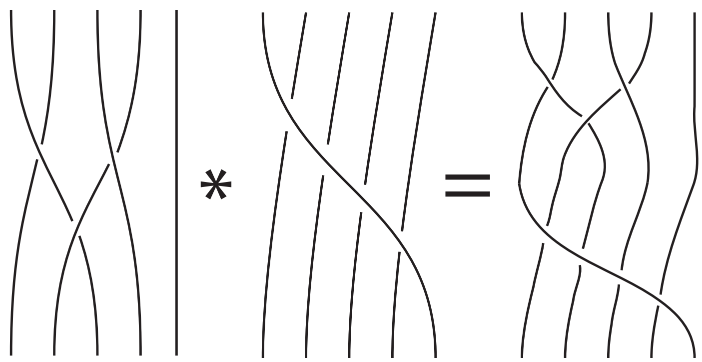

Figure 3.2: The binary operation of the braid group

Artin revised his definition in [2] to add a greater level of rigour to the theory by incorporating real analysis.

Another approach to making the theory more rigorous was introduced by Fox and Neuwirth[13], which takes an algebraic topological approach by constructing the classical braid groups as fundamental groups of a related space. This definition has two advantages over the original: it easily generalizes to define braid groups over an arbitrary topological space, and it enables one to make use of homotopy theory and the functorality of π_1 to prove results.

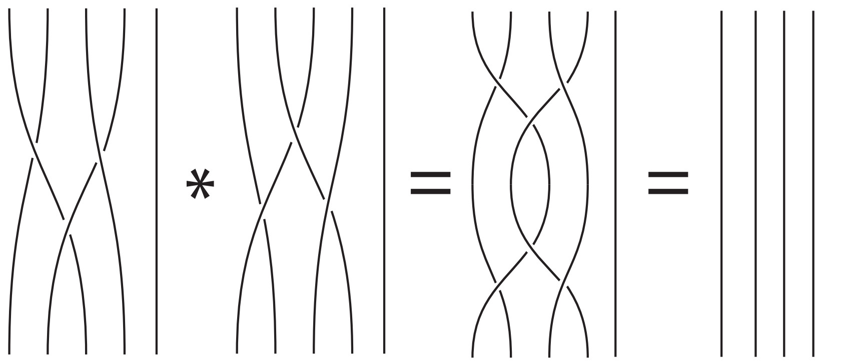

Figure 3.3: The inverse of a braid is its mirror about the horizontal plane

Note that while the definitions below are all rigorous when M is an arbitrary topological space, we will assume throughout this chapter that M is a connected topological m-dimensional manifold, with $m \geq 2$.

Definition 23. *The* ***n-fold configuration space*** *of M is the set*

$$\mathrm{Conf}_n(M) = \{(x_1, \ldots, x_n) \in M^n \mid x_i \neq x_j \textit{ for } i \neq j\}$$

equipped with the subspace topology.

Note that when M is an m-dimensional manifold, $\mathrm{Conf}_n(M)$ is an open submanifold of M^n, and hence has dimension mn.

Fix a basepoint $(x_1^0, \ldots, x_n^0) \in \mathrm{Conf}_n(M)$, and consider $\pi_1(\mathrm{Conf}_n(M))$. Given a path $\gamma : I \to \mathrm{Conf}_n(M)$, there exists a unique path $\gamma_i : I \to M$ whose initial point is x_i^0, defined by projecting $\mathrm{Conf}_n(M)$ onto the i^{th} coordinate. We will refer to this path as the i^{th} strand of γ. For distinct i, j, $\gamma_i(t) \neq \gamma_j(t)$ as $\gamma(t) \in \mathrm{Conf}_n(M)$. We therefore have that the union of the graphs of the γ_i's form a subset of $S^1 \times M$ consisting of n disjoint circles. This object therefore shares a close resemblance to a classical braid. However, under this definition the i^{th} strand of $\gamma \in \pi_1(\mathrm{Conf}_n(M))$ is required to end at x_i^0 for all i, and so we are not constructing all braids. To remedy this, we quotient out by the natural S_n action on $\mathrm{Conf}_n(M)$.

Given $(x_1, \ldots, x_n) \in \mathrm{Conf}_n(M)$, a permutation $\sigma \in S_n$ naturally acts on $\mathrm{Conf}_n(M)$ by $\sigma((x_1, \ldots, x_n)) = (x_{\sigma(1)}, \ldots, x_{\sigma(n)})$. Since this group action is free and S_n is finite, the projection from $\mathrm{Conf}_n(M)$ to the quotient space

$\widehat{\mathrm{Conf}}_n(M) = \mathrm{Conf}_n(M)/S_n$ is a finite cover, and hence $\widehat{\mathrm{Conf}}_n(M)$ is also an mn-dimensional manifold. We can consider $\widehat{\mathrm{Conf}}_n(M)$ to be a topological structure placed upon the subsets of M with cardinality n.

If we now consider $\pi_1(\widehat{\mathrm{Conf}}_n(M))$, we can again define the strands of a loop γ. Place an (arbitrary) linear ordering on the base point $\{x_1^0, \ldots, x_n^0\}$. Using this choice, we have a unique lifting of γ to a path in $\mathrm{Conf}_n(M)$, and we can define the i^{th} strand of γ to be the i^{th} strand of this lifting, again denoting the i^{th} strand of γ as γ_i. Although the label applied to each strand depends on the choice of ordering, throughout this thesis we will view the ordering to be arbitrary but fixed. Again, the union of the graphs of the strands do not intersect.

We now have that the endpoints of γ_i do not necessarily agree. Rather, γ defines a permutation in S_n defined by sending x_i^0 to $\gamma_i(1)$. It is clear that this map $\pi_1(\widehat{\mathrm{Conf}}_n(M)) \to S_n$ is a surjective group morphism.

We can identify the classical braid group with $\pi_1(\widehat{\mathrm{Conf}}_n(\mathbb{R}^2))$ by identifying a loop γ with the union of the graphs of the strands. This identification is made rigorous in Section 1.2 of [5].

The above argument justifies the following definitions.

Definition 24. *The group* $\pi_1(\mathrm{Conf}_n(M))$ *is the* ***pure braid group*** *of* M*. We will denote this group as* $\mathrm{PBr}_n(M)$*.*

Definition 25. *The group* $\pi_1(\widehat{\mathrm{Conf}}_n(M))$ *is the* ***braid group*** *of* M*. We will denote this group as* $\mathrm{Br}_n(M)$*.*

In the above descriptions, we have insinuated that $\mathrm{PBr}_n(\mathbb{R}^2)$ is a subgroup of $\mathrm{Br}_n(\mathbb{R}^2)$. Indeed, more generally, $\mathrm{PBr}_n(M)$ is realized as the kernel of the above mentioned group morphism $\sigma : \mathrm{Br}_n(M) \to S_n$.

Proposition 26. [5] *There exists a short exact sequence*

$$1 \longrightarrow \mathrm{PBr}_n(M) \longrightarrow \mathrm{Br}_n(M) \xrightarrow{\sigma} S_n \longrightarrow 1.$$

Proof. We can define a group morphism $\mathrm{PBr}_n(M) \to \mathrm{Br}_n(M)$ by mapping a loop $\gamma \in \mathrm{PBr}_n(M)$ to $\tilde{\gamma}$, where $\tilde{\gamma}$ is obtained by composing γ with the projection $\mathrm{Conf}_n(M) \to \widehat{\mathrm{Conf}}_n(M)$. Since $\tilde{\gamma}$ lifts to γ, we have that the morphism is an injection into the $\operatorname{Ker}\sigma$. Moreover, it is clear from the definition of σ that loops in $\operatorname{Ker}\sigma$ lift to loops of $\mathrm{Conf}_n(M)$, and thus the group morphism is an isomorphism $\mathrm{PBr}_n(M) \to \operatorname{Ker}\sigma$. □

Note that the permutation associated to a braid provides information on the number of path components of the union of the graphs of the strands of a braid, which is a subset of $S^1 \times M$. Each cycle in the disjoint cycle decomposition of the image of the braid under σ describes how the graphs of the corresponding strands are connected. We therefore have that the number of components of a braid is equal to the number of cycles in the disjoint cycle decomposition of the associated permutation.

Definition 27. *A braid is* ***connected*** *if the union of graphs of strands consists of a single component. Equivalently, a braid is connected if and only if the associated permutation is an n-cycle.*

As an example of the benefit offered by developing the braid groups as fundamental groups, we now present a fibre bundle that establishes relations among the pure braid groups. Before doing so, for each $n \in \mathbb{N} \cup \{0\}$, define a subset $\mathcal{P}_n \subset M$ of cardinality n such that $\mathcal{P}_{n-1} \subset \mathcal{P}_n$ for all n. This induces a natural linear order on $\cup_{n=1}^{\infty} \mathcal{P}_n$.

Lemma 28. [5] *Fix $n > m$. The map $\tau : \mathrm{Conf}_n(M \setminus \mathcal{P}_k) \to \mathrm{Conf}_m(M \setminus \mathcal{P}_k)$ defined by $\tau(x_1, \ldots, x_n) = (x_1, \ldots, x_m)$ is a fibre bundle with fibre homeomorphic to $\mathrm{Conf}_{n-m}(M \setminus \mathcal{P}_{k+m})$.*

Proof. Given $(x_1, \ldots, x_m) \in \mathrm{Conf}_m \setminus \mathcal{P}_k$, the preimage under τ will consist of the points

$$\{x_1, \ldots, x_m, y_1, \ldots, y_{n-m} \mid y_i \neq y_j \text{ for } i \neq j, \ \ y_i \notin \mathcal{P}_k \cup \{x_1, \ldots, x_m\} \text{ for all } i\}$$

This set is clearly homeomorphic to $\mathrm{Conf}_{n-m}(M \setminus \mathcal{P}_{k+m})$, and so τ is a pre-bundle.

To see that this pre-bundle is locally trivial, begin by selecting distinct points $z_1, \ldots, z_m \in M \setminus \mathcal{P}_k$. Pick an orientation-preserving automorphism $\alpha : M \to M$ that fixes $\mathcal{P}_k$ pointwise and maps p_{k+i} to z_i for all $1 \leq i \leq m$. Choose a collection of pairwise disjoint open neighbourhoods U_i of z_i in $M \setminus \mathcal{P}_k$ that are each homeomorphic to an open ball. Next, for each i, define a map $\theta^i : U_i \times \overline{U_i} \to \overline{U_i}$ such that:

- $\theta^i(x, -)$ is an automorphism of $\overline{U_i}$ that fixes $\partial\overline{U_i}$ pointwise
- for all $x \in U_i$, $\theta^i(x, x) = z_i$.

We will also use θ^i_x to denote $\theta^i(x, -)$. Since each θ^i_x fixes $\partial\overline{U_i}$, we can extend θ^i to $U_i \times M \to M$, by demanding that θ^i is trivial outside of U_i. Note that $\tau^{-1}(U_1 \times \cdots \times U_m)$ consists of all points in $\mathrm{Conf}_n(M \setminus \mathcal{P}_k)$ such that the first m points are in the respectively labelled U_i.

Given $x_i \in U_i$ for each $1 \le i \le m$, let

$$\Theta_{x_1,\dots,x_m} = \alpha^{-1} \circ \theta^1_{x_1} \circ \cdots \circ \theta^m_{x_m}.$$

Note that $\Theta_{x_1,\dots,x_m}$ is a homeomorphism from $M \setminus (\mathcal{P}_k \cup \{x_1, \dots, x_m\})$ to $M \setminus \mathcal{P}_{k+m}$ that continuously varies as we vary the x_i's. We can then define a local trivialization about a neighbourhood of $(z_1, \dots, z_m)$ via

$$\varphi : U_1 \times \cdots \times U_m \times \mathrm{Conf}_{n-m}(M \setminus \mathcal{P}_{k+m}) \longrightarrow \tau^{-1}(U_1 \times \cdots \times U_m)$$

where

$$\varphi(x_1, \dots, x_n) = \big(x_1, \dots, x_m, \Theta^{-1}_{x_1,\dots,x_m}(x_{m+1}), \dots, \Theta^{-1}_{x_1,\dots,x_m}(x_n)\big).$$

We therefore have that τ is a fibre bundle. $\square$

Note that the bundle in Lemma 28 does not have a counterpart in the $\widehat{\mathrm{Conf}}_n(M \setminus mathcal P_k)$ setting as, in general, we have no method for choosing a preferred subset of $\{x_1, \dots, x_n\} \in \widehat{\mathrm{Conf}}_n(M \setminus mathcal P_k)$ of size m.

Since $M \setminus \mathcal{P}_{k+m}$ is path connected (under the assumption that $\dim M \ge 2$), applying Theorem 8 provides the following long exact sequence.

Corollary 29. [5] *There exists a long exact sequence*

$$\cdots \longrightarrow \mathrm{PBr}_{n-m}(M \setminus \mathcal{P}_{k+m}) \longrightarrow \mathrm{PBr}_n(M \setminus \mathcal{P}_k) \xrightarrow{\tau_*} \mathrm{PBr}_m(M \setminus \mathcal{P}_k) \longrightarrow 1$$

A specialization of this long exact sequence results in the following theorem.

Theorem 30. [5] *Suppose $\pi_2(M \setminus \mathcal{P}_k) = \pi_3(M \setminus \mathcal{P}_k) = 1$ for all $k \ge 0$. There exists a short exact sequence*

$$1 \longrightarrow \mathrm{PBr}_1(M \setminus \mathcal{P}_{n-1}) \longrightarrow \mathrm{PBr}_n(M) \xrightarrow{\tau_*} \mathrm{PBr}_{n-1}(M) \longrightarrow 1$$

for all $n \ge 2$.

Proof. Fix $l \ge 2$. Specializing Corollary 29 to $k = 0$, $n = l$ and $m = l - 1$, we have a long exact sequence

$$\cdots \to \pi_2(\mathrm{Conf}_{l-1}(M)) \to \mathrm{PBr}_1(M \setminus \mathcal{P}_{l-1}) \to \mathrm{PBr}_l(M) \to \mathrm{PBr}_{l-1}(M) \to 1.$$

It therefore suffices to show that $\pi_2(\mathrm{Conf}_{l-1}(M)) = 1$. To prove this, we will again specialize Corollary 29, this time setting $m = 1$. We then have the following long exact sequence.

$$\cdots \to \pi_3(M \setminus \mathcal{P}_k) \to \pi_2(\mathrm{Conf}_{n-1}(M \setminus \mathcal{P}_{k+1})) \to \\ \to \pi_2(\mathrm{Conf}_n(M \setminus \mathcal{P}_k)) \to \pi_2(M \setminus \mathcal{P}_k) \to \cdots$$

Under our initial assumption, $\pi_3(M \setminus \mathcal{P}_k)$ and $\pi_2(M \setminus \mathcal{P}_k)$ are trivial for all $k \geq 0$ and so for all $k \geq 0$, $n \geq 2$, there exists an isomorphism

$$\pi_2(\mathrm{Conf}_{n-1}(M \setminus \mathcal{P}_{k+1})) \cong \pi_2(\mathrm{Conf}_n(M \setminus \mathcal{P}_k)).$$

We therefore inductively have that

$$\pi_2(\mathrm{Conf}_{l-1}(M)) \cong \pi_2(M \setminus \mathcal{P}_{l-2}).$$

Since $\pi_2(M \setminus \mathcal{P}_{l-2})$ is trivial by the initial assumptions, this completes the proof. □

3.3 The Mapping Class Group

Let $\mathcal{P}_n$ be as in the previous section. Let $\mathrm{Aut}_n(M)$ denote the topological group of orientation-preserving automorphisms of M that map $\mathcal{P}_n$ onto itself and act trivially on ∂M.

Definition 31. *The* ***n-punctured mapping class group*** *of M is defined to be $\pi_0(\mathrm{Aut}_n(M))$. We will denote this group as $\Gamma_n(M)$.*

By definition, $\varphi \in \mathrm{Aut}_n(M)$ restricts to a permutation on $\mathcal{P}_n$. Note that this assignation is a group morphism $\mathrm{Aut}_n(M) \to S_n$ and is locally-constant since $\mathcal{P}_n$ is discrete. We therefore have a group morphism $\Gamma_n(M) \to S_n$.

One subgroup of $\Gamma_n(M)$ that we wish to note is constructed from the set $\mathrm{Aut}_n^0(M)$ of orientation-preserving automorphisms that restrict to the identity on $\mathcal{P}_n$.

Definition 32. *The group $\pi_0(\mathrm{Aut}_n^0(M))$ is the* ***n-punctured pure mapping class group****. We will denote this group as $\mathrm{P}\Gamma_n(M)$.*

Note that $\mathrm{Aut}_0^0(M) = \mathrm{Aut}_0(M)$ and $\mathrm{Aut}_1^0(M) = \mathrm{Aut}_1(M)$ for all M, and thus $\mathrm{P}\Gamma_0(M) = \Gamma_0(M)$ and $\mathrm{P}\Gamma_1(M) = \Gamma_1(M)$ for all M.

More generally, the inclusion $\mathrm{Aut}_n^0(M)) \hookrightarrow \mathrm{Aut}_n(M)$ induces an inclusion $P\Gamma_n(M) \hookrightarrow \Gamma_n(M)$. This is easy to see as the above map $\mathrm{Aut}_n(M) \to S_n$ is locally constant and $\mathrm{Aut}_n^0(M)$ is the kernel of this map. We therefore have that $\mathrm{Aut}_n^0(M)$ is a union of components of $\mathrm{Aut}_n(M)$, which demonstrates that $P\Gamma_n(M)$ is a subgroup of $\Gamma_n(M)$. Indeed, we have proven the following proposition.

Proposition 33. *There exists a short exact sequence*

$$1 \longrightarrow P\Gamma_n(M) \longrightarrow \Gamma_n(M) \longrightarrow S_n \longrightarrow 1.$$

The theory of mapping class groups is closely related to the theory of braid groups. An important connection between the pure braid groups and the pure mapping class groups is given in the following theorem.

Theorem 34. [5] *There exists a fibre bundle* $\pi : \mathrm{Aut}_0^0(M) \to \mathrm{Conf}_n(M)$ *with fibre* $\mathrm{Aut}_n^0(M)$*.*

Proof. Let $\{p_i\}_{i=1}^{\infty}$ be the ordering of $\cup_{i=1}^{\infty}\mathcal{P}_i$ as above. Given $\varphi \in \mathrm{Aut}_0^0(M)$, define $\pi(\varphi) = (\varphi(p_1), \ldots, \varphi(p_n))$. This is clearly a surjective continuous map. Also, the preimage $\pi^{-1}(p_1, \ldots, p_n)$ consists of all automorphisms in $\mathrm{Aut}_0^0(M)$ that restrict to the identity on $\mathcal{P}_n$; this is the subgroup $\mathrm{Aut}_n^0(M)$. Moreover, for any other point $(q_1, \ldots, q_n) \in \mathrm{Conf}_n(M)$, composition using any orientation preserving automorphism that maps q_i to p_i for all $1 \leq i \leq n$ provides an isomorphism between $\pi^{-1}(q_1, \ldots, q_n)$ and $\pi^{-1}(p_1, \ldots, p_n)$. We therefore have that π is a pre-bundle.

Since the fibres have a group structure, according to Exercise 99 of [21], it suffices to find a local section about each point $(z_1, \ldots, z_n) \in \mathrm{Conf}_n(M)$ in order to show that there exists a local trivialization of this pre-bundle. Given such a $z = (z_1, \ldots, z_n)$, define pairwise disjoint open neighbourhoods $U_i \subset M$ about z_i. We then have that $U = U_1 \times \cdots \times U_n$ is an open subset about z. Define a continuous function $U \to \mathrm{Aut}_0^0(M)$ that assigns to each $u = (u_1, \ldots, u_n) \in U$ a function that maps z_i to u_i and is the identity outside of U (such as $(\theta_{u_1}^1 \circ \cdots \circ \theta_{u_n}^n)^{-1}$ in the proof of Lemma 28); this function is a local section over U, completing the proof. □

Using Theorem 8, the long exact sequence corresponding to this fibre bundle contains the following segment.

Corollary 35. [5] *There exists a long exact sequence ending in*

$$\cdots \longrightarrow \pi_1(\mathrm{Aut}_0^0(M)) \xrightarrow{\pi_*} \mathrm{PBr}_n(M) \xrightarrow{d} \mathrm{P}\Gamma_n(M) \xrightarrow{i_*} \mathrm{P}\Gamma_0(M) \longrightarrow 1.$$

The surjection $\mathrm{P}\Gamma_n(M) \to \mathrm{P}\Gamma_0(M)$ is induced by the inclusion morphism $i : \mathrm{Aut}_n^0(M) \hookrightarrow \mathrm{Aut}_0^0(M)$. We can therefore describe this map as simply forgetting that we are fixing $\mathcal{P}_n$. In particular, the kernel of this map consists of the mapping classes that are represented by automorphisms that are isotopic to the identity once we no longer demand that the isotopy preserves $\mathcal{P}_n$. The third map of this long exact sequence describes this kernel as the image of $\mathrm{PBr}_n(M)$.

Intuitively, we can consider the map $d : \mathrm{PBr}_n(M) \to \mathrm{P}\Gamma_n(M)$ as follows. Let $\beta \in \mathrm{PBr}_n(M)$. Fix a representative $(\beta_1, \ldots, \beta_n) : I \to \mathrm{Conf}_n(M)$. We can then construct an isotopy $h_t : M \to M$ such that $h_0 = id$ and $h_t(x_i) = \beta_i(t)$. We then define $d(\beta) = [h_1]$. Since the endpoints of any two isotopies defined from representatives of β are themselves homotopic, this map is well-defined. Since there exist representatives of $d(\beta)$ constructed with a specific isotopy to id_M, we clearly have that $\mathrm{Im}\ d \subset \mathrm{Ker}\, i_*$. Moreover, given $[f] \in \mathrm{Ker}\, i_*$, there exists an isotopy $f_t : I \times M \to M$ such that $f_0 = id_M$ and $f_1 = f$. Since we can define a braid uniquely by its strands, we have a braid β defined by $\beta_i : I \to M$ where $\beta_i(t) = f_t(p_i)$. Moreover, f_t is the desired isotopy that maps β to $[f]$ under d. We therefore have that $\mathrm{Im}\ d = \mathrm{Ker}\, i_*$.

The next step in analyzing this long exact sequence is to determine the image of $\pi_1(\mathrm{Aut}_0^0(M))$ in $\mathrm{PBr}_n(M)$. When M is the genus g Riemann surface Σ_g, the answer is quite nice.

Theorem 36. [5] *Let $g, n \geq 0$. Consider the above long exact sequence*

$$\cdots \longrightarrow \pi_1(\mathrm{Aut}_0^0(\Sigma_g)) \xrightarrow{\pi_*} \mathrm{PBr}_n(\Sigma_g) \xrightarrow{d} \mathrm{P}\Gamma_n(\Sigma_g) \xrightarrow{i_*} \mathrm{P}\Gamma_0(\Sigma_g) \longrightarrow 1.$$

If $g \geq 2$, $g = 1$ and $n \geq 2$, or $g = 0$ and $n \neq 3$, the image of π_* is equal to the centre, $Z(\mathrm{PBr}_n(\Sigma_g))$, of $\mathrm{PBr}_n(\Sigma_g)$. In particular,

$$\mathrm{Ker}\, i_* \cong \mathrm{PBr}_n(\Sigma_g)/Z(\mathrm{PBr}_n(\Sigma_g)).$$

Proof. We will only prove this theorem for $g \geq 2$.

We will begin by proving that $\mathrm{Ker}\, d$ is contained within $Z(\mathrm{PBr}_n(\Sigma_g))$. Let $\beta \in \mathrm{Ker}\, d$. Then there exists an element of $\pi_1(\mathrm{Aut}_0^0(\Sigma_g))$ that maps to β. Say that this element is represented by $h_t : I \times \Sigma_g \to \Sigma_g$. Let γ be another element of $\mathrm{PBr}_n(\Sigma_g)$, and suppose that γ is represented by $(\gamma_1, \ldots, \gamma_n) : I \to$

$\mathrm{Conf}_n(\Sigma_g)$. We can then define a map $I \times I \to \mathrm{Conf}_n(\Sigma_g)$ by sending (s,t) to $(h_t(\gamma_1(s)), \ldots, h_t(\gamma_n(s)))$. The boundary of this square defines a representative of the braid $\beta\gamma\beta^{-1}\gamma^{-1}$, and thus $\beta\gamma\beta^{-1}\gamma^{-1} = 1$. Since this is true for all $\gamma \in \mathrm{PBr}_n(\Sigma_g)$, $\beta \in Z(\mathrm{PBr}_n(\Sigma_g))$.

We will now proceed to show that when $g \geq 2$, $\mathrm{PBr}_n(\Sigma_g)$ is centreless. Since Σ_g is acyclic, we can apply Theorem 30 to obtain the exact sequence

$$1 \longrightarrow \mathrm{PBr}_1(\Sigma_g \setminus \mathcal{P}_{n-1}) \xrightarrow{\partial_*} \mathrm{PBr}_n(\Sigma_g) \xrightarrow{\tau_*} \mathrm{PBr}_{n-1}(\Sigma_g) \longrightarrow 1$$

for all $n \in \mathbb{N}$. Moreover, we have that $\mathrm{PBr}_1(\Sigma_g) = \pi_1(\Sigma_g)$ is centreless. Assume inductively that $\mathrm{PBr}_{n-1}(\Sigma_g)$ is also centreless. Then $\tau_*(Z(\mathrm{PBr}_n(\Sigma_g))) = 1$ and so $Z(\mathrm{PBr}_n(\Sigma_g)) \subseteq \mathrm{Im}\partial_* \cong \mathrm{PBr}_1(\Sigma_g \setminus \mathcal{P}_{n-1})$. Choose a disc D^2 that contains $\mathcal{P}_n$. We can describe $\pi_1(D^2 \setminus \mathcal{P}_n)$ as

$$\langle \gamma_1, \ldots, \gamma_n, z \mid \gamma_1 \cdots \gamma_n = z \rangle$$

where γ_i separates p_i from the rest of $\mathcal{P}_n$ and is properly oriented (z is some loop which partitions $\mathcal{P}_n$ from ∂D^2). We can also describe $\pi_1(\Sigma_g \setminus D^2)$ as

$$\langle \alpha_1, \ldots, \alpha_g, \beta_1, \ldots, \beta_g, z \mid \prod_{i=1}^{g} [\alpha_i, \beta_i] = z \rangle$$

where the α's and β's provide the usual presentation of $\pi_1(\Sigma_g)$ and z is again a loop that is isotopic to the boundary. Using Seifert-van Kampen's Theorem, we obtain a presentation

$$\pi_1(\Sigma_g \setminus \mathcal{P}_n) \cong \langle \alpha_1, \ldots, \alpha_g, \beta_1, \ldots, \beta_g, \gamma_1, \ldots, \gamma_n \mid \prod_{i=1}^{g} [\alpha_i, \beta_i] = \gamma_1 \cdots \gamma_n \rangle.$$

Note that this group is isomorphic to the free group on $2g + n - 1$ generators.

In particular, since $Z(\mathrm{PBr}_n(\Sigma_g))$ is contained within $Z(\pi_1(\Sigma_g \setminus \mathcal{P}_{n-1}))$, which is trivial since $\pi_1(\Sigma_g \setminus \mathcal{P}_{n-1})$ is free with rank larger than 1. We therefore have that $\mathrm{PBr}_n(\Sigma_g)$ is centreless, completing the proof. $\square$

Note that the proof for $g = 0$ can be found in [5], and the proof for $g = 1$ can be found in [4].

The idea behind Theorem 36 extends to the non-pure setting. Consider the map $\pi : \mathrm{Aut}_0(M) \to \widehat{\mathrm{Conf}}_n(M)$ defined by $\pi(f) = \{f(p_1), \ldots, f(p_n)\}$. Each fibre is homeomorphic to $\mathrm{Aut}_n(M)$. Given $\{q_1, \ldots, q_n\} \in \widehat{\mathrm{Conf}}_n(M)$, if we choose pairwise disjoint open neighbourhoods U_i about p_i, then the S_n

action on $U_1 \times \cdots \times U_n$ is trivial, and we can consider $U_1 \times \cdots \times U_n$ as an open neighbourhood of $\{q_1, \ldots, q_n\}$. The local sections defined in the proof of Theorem 34 is therefore still well-defined, and thus π is a fibre bundle.

We therefore have a long exact sequence

$$\cdots \to \pi_1(\mathrm{Aut}_0(M)) \to \mathrm{Br}_n(M) \to \Gamma_n(M) \xrightarrow{i_*} \Gamma_0(M) \to 1. \tag{3.1}$$

Applying reasoning that is similar to the proof of Theorem 36 we thus have the following theorem.

Theorem 37. [5] *Let $g, n \geq 0$. Consider the above long exact sequence*

$$\cdots \longrightarrow \pi_1(\mathrm{Aut}_0(\Sigma_g)) \longrightarrow \mathrm{Br}_n(\Sigma_g) \longrightarrow \Gamma_n(\Sigma_g) \xrightarrow{i_*} \Gamma_0(\Sigma_g) \longrightarrow 1.$$

If $g \geq 2$, $g = 1$ and $n \geq 2$, or $g = 0$ and $n \neq 3$, then

$$\mathrm{Ker}\, i_* \cong \mathrm{Br}_n(\Sigma_g)/Z(\mathrm{Br}_n(\Sigma_g)).$$

This theorem is fully proved in [5].

As a corollary, since $\mathrm{Br}_n(\mathbb{R}^2) \cong \mathrm{Br}_n(D^2)$ (because any isotopy between two embeddings of braid representatives in $\mathbb{R}^2 \times I$ can be assumed to act entirely in $D^2 \times I$) and $\Gamma_0(\mathbb{R}^2) \cong \Gamma_0(D^2)$, we have that $\Gamma_n(\mathbb{R}^2) \cong \Gamma_n(D^2)$.

Additionally, considering the above theorem when $n = 1$ leads to an alternate description of $\Gamma_0(\Sigma_g)$ for surfaces Σ_g. This corollary is referred to in literature as the *Dehn-Nielsen Theorem*.

Corollary 38. *The group $\Gamma_0(\Sigma_g)$ is isomorphic to $\mathrm{Out}(\pi_1(\Sigma_g))$, the outer automorphism group of $\pi_1(\Sigma_g)$.*

The idea behind proving this corollary is the following. Fix a CW-complex structure on Σ_g with one 0-cell and one 2-cell. Any element of $\Gamma_1(\Sigma_g)$ induces an automorphism on $\pi_1(\Sigma_g)$. Conversely, an automorphism of $\pi_1(\Sigma_g)$ induces a map on the 1-skeleton of the CW-complex. Since this also defines a unique way to glue the 2-cell onto the transformed 1-skeleton, we have defined an automorphism of Σ_g that fixes the base point. We therefore have an isomorphism $\Gamma_1(\Sigma_g) \cong \mathrm{Aut}(\pi_1(\Sigma_g))$.

Consider the surjective morphism $\Gamma_1(\Sigma_g) \to \Gamma_0(\Sigma_g)$ described in the above theorem. For a fixed representative f of a mapping class $[f] \in \Gamma_0(\Sigma_g)$, once we choose a homotopy class of paths from $f(x_0)$ to x_0, composition of f with an isotopy that is defined by moving along this path determines an element of $\Gamma_1(\Sigma_g)$. However, the choice of paths can be altered by conjugating the path with some element of $\pi_1(\Sigma_g)$. We therefore have that

$$\Gamma_0(\Sigma_g) \cong \Gamma_1(\Sigma_g)/\pi_1(\Sigma_g) \cong \mathrm{Aut}(\pi_1(\Sigma_g))/\mathrm{Inn}(\pi_1(\Sigma_g)),$$

completing the proof.

3.4 Group Presentations and Group Generators

If one accepts the construction of the classical braid groups of the Euclidean plane and the equivalence between this theory and the modern theory of braid groups, it is clear that $\mathrm{Br}_n(\mathbb{R}^2)$ admits a set of generators that simply interchanges adjacent points. See figure 3.4 for diagrams of these generators for $\mathrm{Br}_4(\mathbb{R}^2)$.

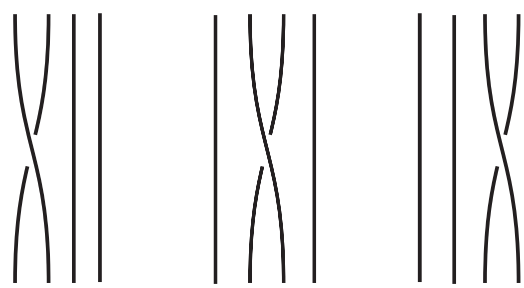

Figure 3.4: The generators of $\mathrm{Br}_4(\mathbb{R}^2)$

As a loop in $\widehat{\mathrm{Conf}}_n(M)$, these generators can be considered as moving two points along a curve such that the union of the curve together with its interior does not contain any marked points. Using these generators, $\mathrm{Br}_n(\mathbb{R}^2)$ admits the following presentation.

Theorem 39. [1] $\mathrm{Br}_n(\mathbb{R}^2)$ *admits the presentation*

$$\left\langle \sigma_1, \ldots, \sigma_{n-1} \,\middle|\, \begin{array}{ll} \sigma_i\sigma_j = \sigma_j\sigma_i & \textit{for} \quad |i-j| \geq 2 \\ \sigma_i\sigma_{i+1}\sigma_i = \sigma_{i+1}\sigma_i\sigma_{i+1} & \textit{for} \quad 1 \leq i \leq n-2 \end{array} \right\rangle.$$

A proof of this is given on page 34 of [5]. Examples of the relations in $\mathrm{Br}_4(\mathbb{R}^2)$ are provided in Figure 3.5.

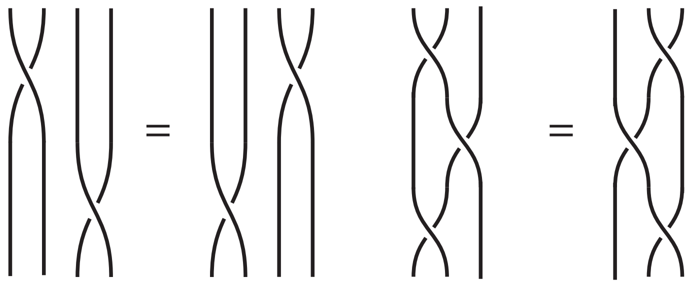

Figure 3.5: Relations in $\mathrm{Br}_4(\mathbb{R}^2)$

The braid group of the sphere admits the following presentation.

Theorem 40. [11] $\mathrm{Br}_n(S^2)$ *admits the presentation*

$$\left\langle \sigma_1, \ldots, \sigma_{n-1} \left| \begin{array}{ll} \sigma_i\sigma_j = \sigma_j\sigma_i & \textit{for} \quad |i-j| \geq 2 \\ \sigma_i\sigma_{i+1}\sigma_i = \sigma_{i+1}\sigma_i\sigma_{i+1} & \textit{for} \quad 1 \leq i \leq n-2 \\ \sigma_1 \cdots \sigma_{n-2}\sigma_{n-1}^2\sigma_{n-2}\cdots\sigma_1 = 1 & \end{array} \right. \right\rangle.$$

This theorem is proved in [11]. Note that we are realizing $\mathrm{Br}_n(S^2)$ as a quotient of $\mathrm{Br}_n(\mathbb{R}^2)$. If we consider the strands of the braids embedded in $S^2 \times S^1$, the $n = 4$ case can be visualized as in Figure 3.6

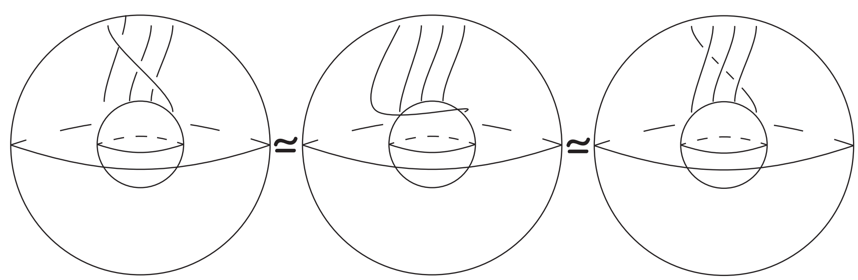

Figure 3.6: $\omega_1\omega_2\omega_3 = \omega_1^{-1}\omega_2^{-1}\omega_3^{-1}$ in $\mathrm{Br}_4(S^2)$

We are now in a position to determine the cardinalities of various subsets of braid groups. In particular, consider braids with the following property.

Definition 41. *Given $\beta \in \mathrm{Br}_n(M)$, identify β with the union of the graphs of the strands of β, which is a disjoint union of circles contained in $S^1 \times M$. Let* *pr* : $S^1 \times M \to M$ *be the canonical projection. A braid is* ***local*** *provided that the product of the projections of the components of* β *under* *pr* *have product* 1 *in* $\pi_1(M)$. *Otherwise,* β *is* ***non-local****.*

Note that all braids of a simply connected manifold are local. We then have the following lemma.

Lemma 42. *Let* $h \geq 0$ *and* $n \geq 1$. *The following sets have the following cardinalities.*

(i)

$$|\mathrm{Br}_n(\Sigma_h)| = \begin{cases} 1 & \text{if} \quad h = 0, \ n = 1, \\ 2 & \text{if} \quad h = 0, \ n = 2, \\ 12 & \text{if} \quad h = 0, \ n = 3, \\ \infty & \text{if} \quad h = 0, \ n \geq 4 \quad \text{or if} \quad h \geq 1. \end{cases}$$

(ii) *Let* $\Delta_{n,h}$ *denote the center of* $\mathrm{Br}_p(\Sigma_h)$. *Then*

$$\Delta_{p,h} \cong \begin{cases} 1 & \text{if} \quad h \geq 2 \quad \text{or if} \quad h = 0, \ n = 1, \\ \mathbb{Z}_2 & \text{if} \quad h = 0, \ n \geq 2, \\ \mathbb{Z} \oplus \mathbb{Z} & \text{if} \quad h = 1. \end{cases}$$

(iii) *Let* $\Lambda_{n,h}$ *be the subset of* $\mathrm{Br}_n(\Sigma_h)$ *consisting of connected local braids. Then*

$$|\Lambda_{n,h}| = \begin{cases} 1 & \text{if} \quad n = 1 \quad \text{or if} \quad h = 0, \ n = 2, \\ 4 & \text{if} \quad h = 0, \ n = 3, \\ \infty & \text{if} \quad h = 0, \ n \geq 4 \quad \text{or if} \quad h \geq 1, \ n \geq 2. \end{cases}$$

(iv) *Let* $q_\Delta : \mathrm{Br}_n(\Sigma_h) \to \mathrm{Br}_n(\Sigma_h)/\Delta_{n,h}$ *be the quotient homomorphism. Then*

$$|q_\Delta(\Lambda_{n,h})| = \begin{cases} 1 & \text{if} \quad n = 1 \quad \text{or if} \quad h = 0, \ n = 2, \\ 2 & \text{if} \quad h = 0, \ n = 3, \\ \infty & \text{if} \quad h = 0, \ n \geq 4 \quad \text{or if} \quad h \geq 1, \ n \geq 2. \end{cases}$$

Proof. The proof of (i) when $h = 0$ follows from Theorem 40. To prove (i) when $h \geq 1$, note that one can inject $\pi_1(\Sigma_h \setminus \mathcal{P}_{n-1})$ into $\mathrm{Br}_n(\Sigma_h)$ by demanding that strands 2 through n are trivial. It then follows that $\mathrm{Br}_n(\Sigma_h)$ is infinite for all positive values of n.

Now consider statement (ii). This statement when $h \geq 2$ was provided in the proofs of Theorems 36 and 37. The proof for $h = 1$ is given in Lemmas 4.2.2 and 4.2.3 of [5]. The proof for $h = 0$ is given in Proposition 4.2 of [29].

Before proving (iii), recall that a braid is connected if and only if the permutation that it maps to under the group morphism of Proposition 26 is an n-cycle.

Thus, since for $h = 0$ all braids are local, precisely $\frac{1}{n}$ of the braids will be both connected and local. When h is arbitrary and $n = 1$ all braids are connected, and so a braid will be local if and only if it is null-homotopic. We therefore have that $|\Lambda_{1,h}| = 1$. This statement for $h \geq 1$ and $n \geq 2$ follows from the fact that we can embed $\mathrm{Br}_n(D^2)$ into $\mathrm{Br}_n(\Sigma_h)$. Consider an injection $D^2 \hookrightarrow \Sigma_h$ such that the points that make up the preferred basepoint of $\mathrm{Br}_n(\Sigma_h)$ are contained within the interior of D^2. It follows from Corollary 4.2 of [29] that, since the complement of D^2 is neither an annulus nor a disk, this embedding leads to an injection $\mathrm{Br}_n(D^2) \hookrightarrow \mathrm{Br}_n(\Sigma_h)$. Since the connected braids in $\mathrm{Br}_n(D^2)$ are connected local braids of $\mathrm{Br}_n(\Sigma_h)$, and $\mathrm{Br}_n(D^2)$ contains an infinite number by Theorem 39, $\mathrm{Br}_n(\Sigma_h)$ contains an infinite number of connected local braids, completing the proof of statement (iii).

Lastly, we will consider statement (iv). The cases where either $n = 1$, $h = 0$ or $h \geq 2$ follow trivially from (iii). When $h = 1$ and $n \geq 2$, the description of $\Delta_{n,1}$ provided in [29] immediately shows that $\mathrm{Br}_n(D^2) \cap \Delta_{n,1} = 1$. Thus, any two distinct braids of T^2 that are contained in $\mathrm{Br}_n(D^2)$ will map to distinct braids under q_Δ, providing an infinite subset of $q_\Delta(\Lambda_{n,1})$. □

Using Theorems 37 and 40, we can then obtain a presentation of $\Gamma_n(S^2)$.

Theorem 43. [5] *If $n \leq 1$, then $\Gamma_n(S^2)$ is trivial. Otherwise $\Gamma_n(S^2)$ admits the presentation*

$$\left\langle \omega_1, \ldots, \omega_{n-1} \;\middle|\; \begin{array}{ll} \omega_i\omega_j = \omega_j\omega_i & \textit{for} \quad |i-j| \geq 2 \\ \omega_i\omega_{i+1}\omega_i = \omega_{i+1}\omega_i\omega_{i+1} & \textit{for} \quad 1 \leq i \leq n-2 \\ \omega_1 \cdots \omega_{n-2}\omega_{n-1}^2\omega_{n-2}\cdots\omega_1 = (\omega_1\cdots\omega_{n-1})^n = 1 & \end{array} \right\rangle.$$

Proof. Specializing Theorem 37 to $n = 1$ provides the exact sequence

$$\pi_1(S^2) \longrightarrow \Gamma_1(S^2) \longrightarrow \Gamma_0(S^2) \longrightarrow \pi_0(S^2)$$

and so $\Gamma_1(S^2) \cong \Gamma_0(S^2)$. However, we also have that $\Gamma_1(S^2) \cong \Gamma_0(\mathbb{R}^2) \cong 1$, which completes the proof for $n = 0, 1$.

For $n \geq 2$, since $\Gamma_0(S^2) \cong 1$, Theorem 37 provides an isomorphism $\Gamma_n(S^2) \cong \mathrm{Br}_n(S^2)/Z(\mathrm{Br}_n(S^2))$. When $n = 2$, Theorem 40 shows that $\mathrm{Br}_2(S^2) \cong \mathbb{Z}_2$ and so $\Gamma_2(S^2) \cong 1$. When $n \geq 3$, the centre of $\mathrm{Br}_n(S^2)$ is the order 2 subgroup generated by $(\sigma_1 \cdots \sigma_{n-1})^n$ (cf. Lemma 4.2.3 of [5]). This provides the given isomorphism. □

Note that a description for the images of the generators of the braid group are given below as 'half-Dehn twists'.

As a consequence of theorem 43, it is easy to verify that the map $\rho : \Gamma_4(S^2) \to GL_2(\mathbb{Z})/\pm I$ defined by sending

$$\omega_1 \mapsto \pm \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \qquad \omega_2 \mapsto \pm \begin{bmatrix} 1 & 0 \\ -1 & 1 \end{bmatrix} \qquad \omega_3 \mapsto \pm \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \tag{3.2}$$

is a group morphism. In particular, the absolute value of the trace of the image of a mapping class under ρ is a well-defined invariant of the conjugacy class of that mapping class.

For $g \geq 1$, we can describe a set of generators of $\Gamma_n(\Sigma_g)$ as follows. Given a simple closed curve γ on Σ_g, let A be an orientation-preserving embedding of the annulus $I \times S^1$ into Σ_g with $A(\frac{1}{2}, t) = \gamma(t)$. Define the Dehn twist about γ to be the automorphism that is trivial outside of an annulus surrounding γ, and acts upon A via the map $(t, z) \mapsto (t, e^{2\pi i t} \cdot z)$ (viewing S^1 as $U(1)$). This can action can be seen in Figure 3.7.

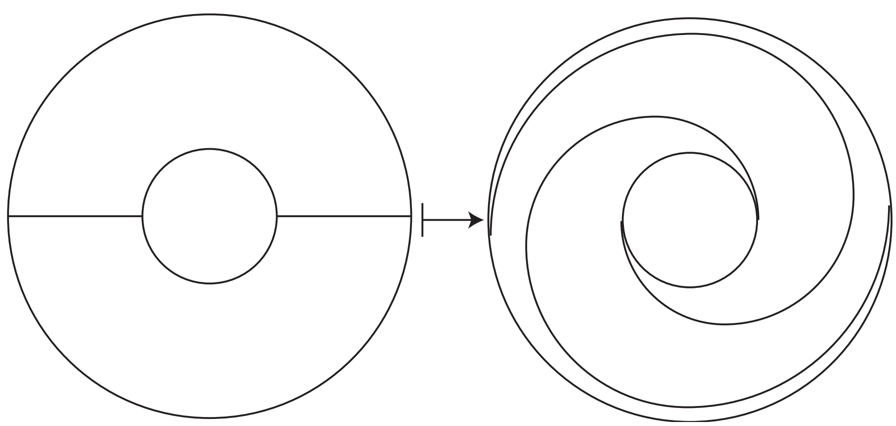

Figure 3.7: A Dehn Twist

Note that if the curve γ is null-homotopic and contracts along a disk whose intersection with $\mathcal{P}_n$ has cardinality less than 2, then the Dehn twist along γ represents the trivial mapping class.

Next consider a simple path between $\lambda : I \to \Sigma_g$ between two points of $\mathcal{P}_n$ such that $\lambda((0,1)) \cap \mathcal{P}_n = \emptyset$. Viewing D^2 as the points in $\mathbb{C}$ with norm less than or equal to 1, consider an orientation-preserving embedding $\varphi : D^2 \hookrightarrow \Sigma_g$

such that $\lambda(t) = \varphi(t - \frac{1}{2})$ and $\text{Im}\varphi \cap \mathcal{P}_n = \{\lambda(0), \lambda(1)\}$. We can then define a half-Dehn twist about λ to be the automorphism of Σ_g that is trivial outside of D^2 and acts on D^2 by the map $z \mapsto e^{2\pi iz} \cdot z$. This action can be seen in Figure 3.8

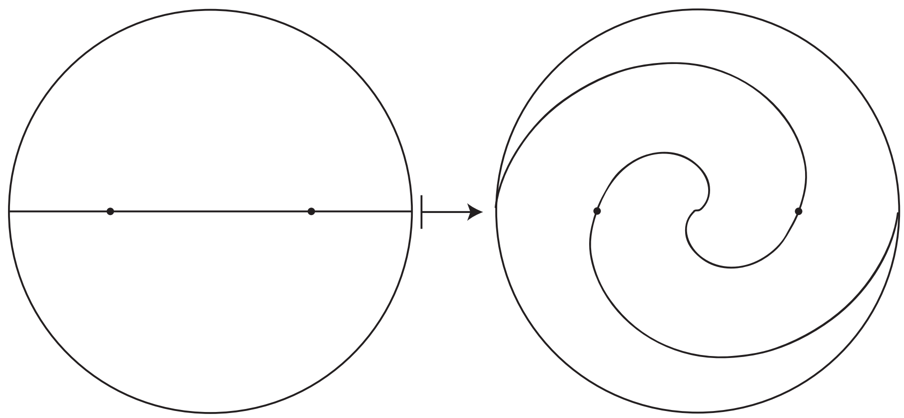

Figure 3.8: A Half-Dehn Twist

According to Corollary 2.11 of [23], we can describe a generating set of $\Gamma_n(\Sigma_g)$ in terms of Dehn twists and half-Dehn twists.

Theorem 44. [23] *Let $g \geq 1$ and $n \geq 0$. $\Gamma_n(\Sigma_g)$ is generated by Dehn twists about curves $a_0, a_1, b_1, \ldots b_{2g-1}, c$ and half-Dehn twists about $\tau_1, \ldots, \tau_{g-1}$, where these curves are depicted in Figure* 3.9*.*

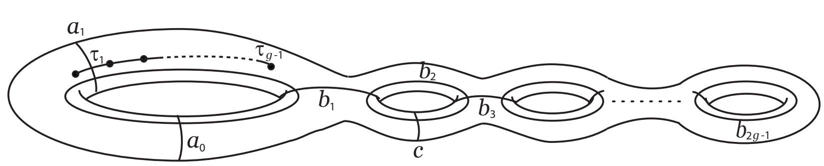

Figure 3.9: Generators of $\Gamma_n(\Sigma_g)$

Note that the half-Dehn twist permutes $\lambda(0)$ and $\lambda(1)$. In [23], the half-Dehn twists are referred to as 'braid twists' because they satisfy the braid relations; given an ordering of $\mathcal{P}_n$, we then have that $\omega_i\omega_j = \omega_j\omega_i$ for $|i - j| > 1$ and

$\omega_i\omega_{i+1}\omega_i = \omega_{i+1}\omega_i\omega_{i+1}$, where ω_i denotes the half-Dehn twist about p_i and p_{i+1}.

3.5 The ζ-Mapping Class Group

Throughout this section, we will consider $\pi : \Sigma_k \to \Sigma_h$ to be a cyclic n-fold branched cover with branch locus $X = \{x_1, \dots, x_p\}$.

We wish to develop a criterion for determining whether a closed curve in $\Sigma_h \setminus X$ will lift to a closed curve in the covering space.

Alternatively, this problem can be reinterpreted as follows. Fix a decomposition of $\Sigma_k \setminus \pi^{-1}(X)$ into n sheets. Since the branched cover is cyclic, the monodromy representation is a group morphism $\pi_1(\Sigma_h \setminus X) \to \mathbb{Z}_n$. An element of $\pi_1(\Sigma_h \setminus X)$ will lift to a closed curve if and only if it induces the trivial automorphism under this representation.

At each branch point $x_i \in X$, choose a simple closed curve $\alpha_i \in \pi_1(\Sigma_h \setminus X)$ that partitions x_i from the other x_j's. These choices should be made in a uniform manner such that they all provide the same action on $\mathbb{Z}_n$. Choose a disk $D^2 \subset \Sigma_h$ that contains X. Applying Seifert-van Kampen's Theorem along the sets $\Sigma_h \setminus D^2$ and $D^2 \setminus X$ provides a presentation

$$\pi_1(\Sigma_h \setminus X) \cong \Big\langle \alpha_1, \dots, \alpha_p, \gamma_1, \dots, \gamma_h, \delta_1, \dots, \delta_h \mid \alpha_1 \cdots \alpha_p = \prod_{i=1}^{h} [\gamma_i, \delta_i] \Big\rangle.$$

Suppose that one has a curve γ in $\Sigma_h \setminus X$ that lifts to a closed curve. Outside of D^2, if one were to remove a segment of γ and replace it with any other curve that does not enter D^2, the resulting curve would also lift to a closed curve. Thus, the only part of the curve that matters for the above consideration is that which lies in D^2. Note that one can construct a new (topological) branched covering as follows. Delete $\Sigma_h \setminus D^2$ as well as $\pi^{-1}(\Sigma_h \setminus D^2)$. The preimage of the ∂D^2 is n disjoint circles, and so we can cap off all $n+1$ circles with discs, and choose an orientation preserving map from the covering discs to the base disc that respects the map along the boundary. We have therefore constructed a branched covering with branch locus X whose base space is S^2 and whose covering space has genus $k - nh$. Denote this branched covering as $\pi' : \Sigma_{k-nh} \to S^2$.

Choosing a base point $z_0 \in D^2 \setminus X$, there exists a group morphism $\pi_1(\Sigma_h \setminus$

$X, z_0) \to \pi_1(S^2 \setminus X, z_0)$ defined by deleting the arcs of the curve that lie outside of D^2, mapping the remaining arcs to the $D^2 \setminus X$ contained in S^2, and completing the missing arcs by *any* arc that is contained in the complement disc. Using the above mentioned presentations, we can define this morphism by sending the α_i's to themselves, and sending the γ_i's and δ_i's to 1. Using the above argument, a curve will lift to a closed curve if and only if its image under this map also lifts to a closed curve.

Considering the special case where the base space is S^2, the fundamental group is generated by $\alpha_1, \ldots, \alpha_{p-1}$ (since $\alpha_p = (\alpha_1 \ldots \alpha_{p-1})^{-1}$), and by the above imposed conditions on the covering space, each α_i maps to the same element under the monodromy representation. Thus, a word in the α_i's will induce the trivial automorphism if and only if n divides the number of α_i's (counted with sign).

This reasoning motivates the 'α-length' group morphism $\ell_\alpha : \pi_1(\Sigma_h \setminus X) \to \mathbb{Z}$, which we define on the generators of $\pi_1(\Sigma_h \setminus X)$ as:

$$\ell_\alpha(\xi) = \begin{cases} 1 & \text{if } \xi = \alpha_1, \ldots, \alpha_{p-1}, \\ 1-p & \text{if } \xi = \alpha_p, \\ 0 & \text{if } \xi = \gamma_1, \ldots, \gamma_h, \delta_1, \ldots, \delta_h. \end{cases} \tag{3.3}$$

Lemma 45. *A closed curve representing $\xi \in \pi_1(\Sigma_h \setminus X)$ will lift to a closed curve in Σ_k under the cyclic branched covering if and only if n divides $\ell_\alpha(\xi)$.*

Proof. Define subgroups W and K of $\pi_1(\Sigma_h \setminus X)$ as follows. Let K consist of all curves in $\pi_1(\Sigma_h \setminus X)$ that lift to a closed curve in $\pi_1(\Sigma_k \setminus \pi^{-1}(X))$. Let W consist of the preimage of $n\mathbb{Z}$ under ℓ_α. Both W and K can be realized as kernels of morphisms: W is the kernel of the composition of ℓ_α with $\mathbb{Z} \to \mathbb{Z}_n$ and K is the kernel of the monodromy representation $\pi_1(\Sigma_h \setminus X) \to \mathrm{Aut}(F)$. Thus, both W and K are normal subgroups, and both of the quotient groups are isomorphic to $\mathbb{Z}_n$. Moreover, n-fold products of the α_i's provide trivial monodromy, and so $W \leq K$.

Thus, a single use of the third isomorphism theorem results in

$$\mathbb{Z}_n \cong \pi_1(\Sigma_h \setminus X)/K \cong \frac{\pi_1(\Sigma_h \setminus X)/W}{K/W} \cong \frac{\mathbb{Z}_n}{K/W}$$

and so $K = W$, completing the proof. □

Lemma 46. *Let $\varphi : \Sigma_h \setminus X \to \Sigma_h \setminus X$ be an orientation-preserving automorphism. For any $\xi \in \pi_1(\Sigma_h \setminus X)$, $\ell_\alpha(\xi) = \ell_\alpha(\varphi_* \xi)$.*

Proof. First consider $h = 0$. Consider the presentation of $\Gamma_n(S^2)$ provided in Theorem 43. Direct computation reveals that for each $i = 1, \ldots, p-1$

$$(\omega_i)_* \alpha_j = \begin{cases} \alpha_j & \text{if} \quad j \neq i, i+1, \\ \alpha_{i+1} & \text{if} \quad j = i, \\ \alpha_{i+1}\alpha_i\alpha_{i+1}^{-1} & \text{if} \quad j = i+1. \end{cases}$$

We therefore have that the α-length of an element of $\Gamma_n(S^2)$ is invariant under the natural $\Gamma_n(S^2)$ action.

Now consider $h \geq 1$. Consider the generators of $\Gamma_n(\Sigma_h)$ provided in Theorem 44. It is clear that the actions of a_0, the b_i's and c map the α_i's onto themselves. The actions of a_1 and the τ_i's act upon the α_i's according to:

$$\begin{aligned} (\tau_i)_* \alpha_j &= \begin{cases} \alpha_j & \text{if} \quad j \neq i, i+1, \\ \alpha_i^{-1}\alpha_{i+1}\alpha_i & \text{if} \quad j = i, \\ \alpha_i & \text{if} \quad j = i+1. \end{cases} \\ (a_1)_* \alpha_j &= \begin{cases} \alpha_j & \text{if} \quad j \neq 1, \\ a_1\, \alpha_1 a_1^{-1} & \text{if} \quad j = 1. \end{cases} \end{aligned}$$

Again, the α-length of an element of $\pi_1(\Sigma_h \backslash X)$ is invariant under the $\Gamma_n(\Sigma_h)$ action, completing the proof. □

As corollary of the previous Lemmas 45 and 46, a curve in $\Sigma_h \setminus X$ lifting to a closed curve is preserved under orientation-preserving automorphisms of $\Sigma_h \setminus X$.

Definition 47. *For a fixed orientation-preserving automorphism $\zeta : \Sigma_k \to \Sigma_k$, let $\mathrm{Aut}^\zeta(\Sigma_k)$ denote the space of all orientation-preserving automorphisms of Σ_k that commute with ζ. Define the ζ-mapping class group to be $\Gamma^\zeta(\Sigma_k) = \pi_0(\mathrm{Aut}^\zeta(\Sigma_k))$.*

This is a generalization of the hyper-elliptic mapping class group introduced in [3]. We are primarily interested in this group when the action of ζ on Σ_k induces a branched cover $\Sigma_k \to \Sigma_k/\langle\zeta\rangle \cong \Sigma_h$. In this case, the mapping class group of Σ_h can be realized a quotient of $\Gamma^\zeta(\Sigma_k)$.

Theorem 48. *Let $\zeta : \Sigma_k \to \Sigma_k$ be an orientation-preserving automorphism that induces an n-fold cyclic branched covering $\pi : \Sigma_k \to \Sigma_h$ with branch locus $X = \{x_1, \ldots, x_p\}$. There exists an isomorphism $\Phi : \Gamma^\zeta(\Sigma_k)/\langle\zeta\rangle \to \Gamma_p(\Sigma_h)$.*

Proof. Given $\varphi \in \mathrm{Aut}^{\zeta}(\Sigma_k)$, define $\Phi(\varphi)$ using the following construction. Since the homeomorphisms in $\mathrm{Aut}^{\zeta}(\Sigma_k)$ are precisely the fibre-preserving automorphisms of π, $\pi\varphi\pi^{-1}$ is a well-defined set map of Σ_h. Indeed, this map is continuous and admits an inverse $\pi\varphi^{-1}\pi^{-1}$, and so is an automorphism of Σ_h. This is a group morphism.

Since ζ maps each fibre onto itself, the map $\pi\zeta\pi^{-1} = id_{\Sigma_h}$. We therefore have a morphism $\mathrm{Aut}^{\zeta}(\Sigma_k)/\langle\zeta\rangle \to \mathrm{Aut}(\Sigma_h)$.

The morphism Φ is an isomorphism that admits the following inverse. Let $f : \Sigma_h \setminus X \to \Sigma_h \setminus X$ be an orientation preserving automorphism. Choose a point $z_0 \in \Sigma_k \setminus \pi^{-1}(X)$. One must choose a point $f_*(z_0) \in \pi^{-1} f \pi(z_0) \subset \Sigma_k \setminus \pi^{-1}(X)$, which will be the image of z_0 in the resulting automorphism of Σ_k. For any other $y \in \Sigma_k \setminus \pi^{-1}(X)$, there exists a path λ from z_0 to y. The path $f\pi\lambda$ lifts to a unique path in Σ_k starting at $f_*(z_0)$. Denote the endpoint of this path by $f_*(y)$.

To see that $f_*(y)$ is well-defined, suppose we have two paths λ_1, λ_2, both starting at z_0 and ending at y. The path $\pi(\lambda_1 \sharp (-\lambda_2))$ lifts to a closed path, and thus by Lemmas 45 and 46, so does $f\pi(\lambda_1 \sharp (-\lambda_2))$. A necessary result of this is that the lifts of $f\pi\lambda_1$ and $f\pi\lambda_2$ starting at $f_*(z_0)$ both have the same endpoint, and so $f_*(y)$ is well defined.

Now f_* is uniquely constructed up to a fixed choice of element in the fibre $\pi^{-1} f \pi(z_0)$. Altering this choice will define a new automorphism of Σ_k that differs from the original by a composition with ζ. We therefore have a well-defined map $\Gamma_p(\Sigma_h) \to \Gamma^{\zeta}(\Sigma_k)/\langle\zeta\rangle$.

This map is easily seen to be the inverse of Φ. □

Chapter 4

Symplectic Topology

4.1 Introduction

The purpose of this chapter is twofold.

First, we will introduce the adjunction formula, which relates the genus of a symplectic surface of a 4-dimensional symplectic manifold to the self-intersection of the surface. This is presented in Section 4.4

The second purpose of this chapter is to provide a special case of Gompf's construction that is essential in the construction provided in Chapter 5. In [14], Gompf provides a method for constructing a symplectic manifold from an initial symplectic manifold by gluing together two disjoint closed codimension 2 symplectic submanifolds that are symplectomorphic and satisfy one submanifold have inverse Euler classes. We consider the special case where our initial manifold has two components with two codimension 2 symplectomorphic submanifolds (one in each component) that have trivial Euler class. This special case is developed in Section 4.5

Sections 4.2 and 4.3 provide initial definitions.

4.2 Symplectic Vector Spaces

A symplectic form on a smooth manifold M is a 2-form ω such that:

- ω is closed (i.e. $d\omega = 0$ where d is the exterior differential),

- ω_p is non-degenerate for all $p \in M$.

A symplectic manifold, then, is a manifold M together with such a 2-form.

We will begin studying symplectic forms locally. A *symplectic vector space* is a vector space V together with a non-degenerate skew-symmetric bilinear form ω on V. Such a vector space is the local model for the tangent space of a symplectic manifold.

In anticipation towards studying submanifolds of a symplectic manifold, we will consider subspaces of a symplectic vector space. Given a subspace W of a symplectic vector space (V, ω), the *ω-perpendicular subspace* of W is the subspace

$$W^\omega = \{v \in V \mid \omega(v, w) = 0 \text{ for all } w \in W\}.$$

Unlike the non-degenerate *symmetric* bilinear forms, a subspace and its ω-perpendicular complement need not intersect trivially. Indeed, these two subspaces will intersect trivially if and only if they satisfy the below mentioned property of being symplectic subspaces.

A subspace W of V is a *symplectic subspace* if the restriction of ω to $W \times W$ is symplectic. Note that the only criterion that one needs to check is that $\omega|_{W\times W}$ is non-degenerate.

These two notions are related in the following lemma.

Lemma 49. *Let (V, ω) be a symplectic vector space. Let W be a subspace of V. The following are equivalent.*

- *W is a symplectic subspace of V.*
- $W \cap W^\omega = \{0\}$.
- $V = W \oplus W^\omega$.

Proof. Using the above definition, W is symplectic if and only if for each $w \in W \setminus \{0\}$ there exists a $u \in W$ such that $\omega(w, u) \neq 0$. This is true if and only if $w \notin W^\omega$, showing that the first and second statements are equivalent.

To prove that the second and third statements are equivalent, it suffices to show that $W + W^\omega = V$ for all subspaces W of V. Let $\{v_1, \dots, v_n\}$ be a basis of V such that $\{v_1, \dots, v_k\}$ is a basis of W. Let $\{v_1^*, \dots, v_n^*\}$ be the dual basis of V^*. Since ω is non-degenerate, we can define a vector space isomorphism $V \to V^*$ by sending x to $\omega(x, -)$. Note that the image of W^ω contains $\{v_{k+1}^*, \dots v_n^*\}$, and so W^ω contains $\{v_{k+1}, \dots, v_n\}$. We therefore have that $W + W^\omega = V$, completing the proof. $\square$

There exist two immediate obstructions to the existence of a symplectic form on M. First, M must have even dimension. This can be proved locally by showing that any symplectic vector space must have even dimension. Since any matrix representation of a symplectic form ω on V is skew-symmetric and non-degenerate, the determinant of ω satisfies $\det(\omega) = \det(-\omega^{tr}) \neq 0$ (where ω^{tr} is the transpose of ω). Thus,

$$\begin{aligned} \det(\omega^2) &= \det(\omega) \cdot \det(-\omega^{tr}) \\ &= \det(\omega) \cdot \det(-I^n) \cdot \det(\omega^{tr}) \\ &= (-1)^n \det(\omega^2) \end{aligned}$$

where n is the dimension of V. We therefore have that every symplectic manifold has even dimension.

Additionally, for M to admit a symplectic form, M must be orientable: if we consider M to be a $2n$-dimensional manifold, the n-fold wedge product of a symplectic form ω on M is a volume form since ω is non-degenerate. In particular, a symplectic manifold has a preferred volume form.

Lemma 50. *Let (M, ω) be a $2n$-dimensional symplectic manifold. The $2n$-form $\omega^n = \underbrace{\omega \wedge \ldots \wedge \omega}_{n}$ is a volume form.*

Proof. It suffices to show that for each $p \in M$, w_p^n does not vanish. Let W be a $2n$-dimensional symplectic vector space. We claim that there exists a decomposition of W into a direct sum of n 2-dimensional symplectic subspaces $W = \oplus_{i=1}^n W_i$ where $W_i^\omega = W_1 \oplus \cdots \oplus W_{i-1} \oplus W_{i+1} \oplus \cdots \oplus W_n$. We will proceed inductively on n.

The case $n = 1$ is trivial since ω is non-degenerate.

Suppose that the lemma is true for all symplectic vector spaces of dimension less than $2n$. Then, since ω is non-degenerate, there exist v_1, $v_2 \in T_pM$ such that $\omega(v_1, v_2) = 1$. Let W_1 be the span of v_1 and v_2. Since W_1 is a symplectic vector space, W_1^ω is disjoint from W_1 and also symplectic (since $(W^\omega)^\omega = W$). Therefore, by the induction hypothesis, $W_1^\omega = \oplus_{i=2}^n W_i$. For $i \geq 2$, W_i^ω contains W_1 and $W_2 \oplus \cdots \oplus W_{i-1} \oplus W_{i+1} \oplus \cdots \oplus W_n$, and hence $W_i^\omega \subseteq W_1 \oplus \cdots \oplus W_{i-1} \oplus W_{i+1} \oplus \cdots \oplus W_n$. Moreover, since the dimensions agree, we in fact have equality, completing the claim.

We can therefore find a basis $v_1, \ldots, v_{2n}$ of W such that $\omega(v_{2i-1}, v_{2i}) = 1$ for all $1 \leq i \leq n$ and $\omega(v_{2i-1}, v_j) = 0$ for $j \neq 2i$. Then,

$$\begin{aligned}
\omega_p^n(v_1, \ldots, v_{2n}) &= \frac{1}{2^n} \sum_{\sigma \in S_{2n}} \operatorname{sgn}(\sigma) \cdot \omega(v_{\sigma(1)}, v_{\sigma(2)}) \cdots \omega(v_{\sigma(2n-1)}, v_{\sigma(2n)}) \\
&= \frac{1}{2^n} \sum_{\sigma \in S_n} 2^n \cdot \omega(v_{\sigma(1)}, v_{\sigma(2)}) \cdots \omega(v_{\sigma(2n-1)}, v_{\sigma(2n)}) \\
&= \frac{1}{2^n} \cdot n! \cdot 2^n \\
&= n!
\end{aligned}$$

where the second summation is over the embedding of S_n in S_{2n} of permutations σ such that $\sigma(2i) = \sigma(2i-1) + 1$ for all $1 \leq i \leq n$. Thus, ω^n does not vanish on W. In particular, w^n does not vanish at any $p \in M$ and so w^n is a volume form, completing the proof. □

When M is a closed manifold, there exists an additional obstruction to the existence of a symplectic form on M. Since ω is closed, ω represents some cohomology element $[\omega] \in H^2(M, \mathbb{R})$. Moreover, the corresponding volume form ω^n is represented by $[\omega^n] = [\omega] \cup \cdots \cup [\omega]$ which generates $H^{2n}(M)$. Thus, there must exist an element $\alpha \in H^2(M, \mathbb{R})$ such that $\alpha^n \neq 0$. In particular, $H^2(M, \mathbb{R})$ cannot be trivial.

Note that when M is 2-dimensional, symplectic forms on M are precisely area forms on M. Thus, being oriented is equivalent to being symplectic. We therefore have that the only closed symplectic 2-manifolds are the genus g surfaces Σ_g equipped with an area form. Since the area forms under the compact-open topology consist of two path components (each consisting of forms that prescribe a certain orientation), there are essentially only two symplectic forms on each of these manifolds.

One method for creating symplectic manifolds that will be used later is to take the direct product of two symplectic manifolds. Given symplectic manifolds (M, ω_M) and (N, ω_N), consider the natural projections $\pi_M : M \times N \to M$ and $\pi_N : M \times N \to N$. We can then define a natural symplectic form ω on $M \times N$ by $\omega = \pi_M^* \omega_M + \pi_N^* \omega_N$. We refer to ω as the *product form* on $M \times N$. Both $\pi_M^* \omega_M$ and $\pi_N^* \omega_N$ are closed on $M \times N$, and thus so is ω. Moreover, since $T_{(x,y)}(M \times N)$ naturally splits as $T_x(M) \times T_y(N)$, for any $v \in T_{(x,y)}(M \times N)$ there exists a tangent vector w such that either $\pi_M^* \omega(v, w) \neq 0$ or $\pi_N^* \omega(v, w) \neq 0$. Moreover, we can chose w such that precisely one of these values is non-zero, which shows that ω is non-degenerate and hence symplectic.

4.3 Symplectic Submanifolds

A submanifold of M is a manifold X together with an embedding $i : X \to M$. We may identify such a submanifold with the image of i equipped with the subspace topology and smooth structure.

Now suppose that (M, ω_M) and (X, ω_X) are symplectic manifolds, and let $i : X \to M$ be an embedding . We say that X is a *symplectic submanifold* of M provided that $i^*\omega_M = \omega_X$. Since the symplectic form on X is determined by the symplectic form on M and i, we will discuss symplectic submanifolds without explicitly identifying their symplectic forms. When doing so, we need only check that $i^*\omega_M$ is symplectic. Note that $i^*\omega_M$ is necessarily closed, and so we are only required to verify that $i^*\omega_M$ is non-degenerate. This is equivalent to verifying that the inclusion of T_xX in T_xM defined by i is a symplectic vector space of T_xM for each $x \in X$.

As an example of a symplectic submanifold, consider the direct product of symplectic manifolds (M, ω_M) and (N, ω_N) equipped with the product form. We then have that the restriction of $\pi_M^*\omega_M + \pi_N^*\omega_N$ to a submanifold $M \times \{pt\}$ is $\pi_M^*\omega_M$, and so $M \times \{pt\}$ is symplectic.

4.4 The First Chern Class of a Symplectic Manifold

The purpose of this section is to better understand symplectic submanifolds. To do so, we will associate symplectic vector bundles to a symplectic submanifold, and study these symplectic bundles. In an example of similarity between complex and symplectic geometry, we can extend Chern classes to the symplectic case. Doing so allows us to define the adjunction formula, which relates the genus of a symplectic submanifold with the self-intersection of the submanifold and the intersection of the submanifold with the canonical class. As in Chapter 2, we will assume that all vector bundles have fixed base space M unless stated otherwise.

Definition 51. *A symplectic vector bundle is a vector bundle such that the fibres are equipped with a smoothly varying symplectic bilinear form.*

Such a bundle can be given a trivialization with structure group $Sp(2n)$, the group of symplectomorphisms from $\mathbb{R}^{2n}$ onto itself, by proceeding in the same manner that provided $O(n)$ as a structure group when given an orthogonal structure.

The easiest examples of symplectic vector bundles are the tangent bundles of a symplectic manifold; a symplectic form on a manifold M satisfies the definition of a smoothly varying symplectic bilinear form on TM. Moreover, given a symplectic submanifold X of M, it follows from Lemma 49 that the normal bundle NX is also symplectic.

When developing invariants of symplectic vector bundles, it is convenient to 'forget' a certain amount of structure: given a symplectic vector bundle, we can consider the underlying *complex vector bundle*.

Definition 52. *A complex vector bundle is a vector bundle E_π together with a vector bundle automorphism $J : E \to E$ that satisfies $J^2((x, y)) = (x, -y)$.*

The idea behind such an automorphism is that J mimics multiplication by $i \in \mathbb{C}$. Indeed, as will be seen in Proposition 53, given a complex vector bundle we actually can organize the fibres as complex vector spaces. Note that we are distinguishing the notion of a complex vector bundle from that of a holomorphic vector bundle since we have no notion of whether a map is holomorphic.

Proposition 53. *Let E_π be a complex vector bundle. Then $\pi^{-1}(x)$ has even dimension. Moreover, we can smoothly equip the fibres with a complex vector space structure.*

Proof. Let V be an n-dimensional vector space, $n \neq 0$, that admits an automorphism $J : V \to V$ satisfying $J^2 = -id$. We will inductively construct a basis of V that has even cardinality. Let $v \in V$, $v \neq 0$. Then $J(v) \notin \text{span}\{v\}$, for otherwise $v = a \cdot J(v)$ for some $a \in \mathbb{R}^*$, implying that

$$J(v) = J(a \cdot J(v)) = -a \cdot v = -a^2 \cdot J(v),$$

which is a contradiction. We therefore have a 2-dimensional subspace $\text{span}\{v, J(v)\} \leq V$. Now suppose there exists a $2k$-dimensional subspace

$W \leq V$ spanned by $\{v_1, \dots, v_k, J(v_1), \dots J(v_k)\}$. Then, for any $v \in V$,

$$\begin{aligned} v \in W \quad &\text{iff} \quad v = a_1 v_1 + \cdots + a_k v_k + J(a_{k+1} v_1 + \cdots + a_{2k} v_k) \\ &\text{iff} \quad J(v) = J(a_1 v_1 + \cdots + a_k v_k) - (a_{k+1} v_1 + \cdots + a_{2k} v_k) \\ &\text{iff} \quad J(v) \in W. \end{aligned}$$

Thus, if $W \neq V$, we can extend W to a $2(k+1)$-dimensional subspace $W \oplus \{v, J(v)\}$ for $v \in V \setminus W$.

We therefore have that V is even dimensional. Moreover, the above basis provides a complex vector space structure on V by defining $(a+ib)v = av+bJ(v)$. When considering a complex vector bundle E_π, we can smoothly associate this complex vector space structure to the fibres since we can extend this basis to spanning frames over the preimages $\pi_\alpha^{-1}(U_\alpha)$ of any locally finite local trivialization of the bundle $\{h_\alpha : \pi^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^n\}_{\alpha \in \Lambda}$ and glue the resulting complex structures together using a partition of unity subordinate to the cover. □

Since a complex vector bundle has a smoothly varying complex structure on the fibres, the structure group is naturally $GL_n(\mathbb{C})$. The same reasoning that reduced the structure group of a real vector bundle from $GL_n(\mathbb{R})$ to $O(n)$ works to reduce the structure group of a complex vector bundle from $GL_n(\mathbb{C})$ to $U(n)$. Moreover, any bundle that admits $U(n)$ as its structure group with the canonical representation presupposes a complex structure. We could therefore Alternately define the complex vector bundles as those equipped with a $U(n)$ trivialization.

There is a natural bijective correspondence between symplectic and complex vector bundles. One method for seeing this is to prove that the classifying spaces $BU(n)$ and $BSp(2n)$ are homotopy equivalent. A more hands-on approach to this correspondence is provided below.

Definition 54. *Given a symplectic bundle (E_π, ω), we say that a complex structure J on E_π is **compatible** if $\omega_x(Jv, Jw) = \omega_x(v, w)$ and $\omega_x(v, J(v)) > 0$ for all $v, w \in \pi^{-1}(x)$.*

If we consider J as equipping the fibres with a complex structure, the first criterion for compatibility states that ω can be extended to be $\mathbb{C}$-bilinear. The second criterion simply fixes a preferred real orientation on each complex line.

For a fixed symplectic structure ω, consider the space of all compatible complex structures.

Theorem 55. [26] *Let* (E_π, ω) *be a symplectic vector bundle. Let* $\mathcal{J}(E_\pi, \omega)$ *be the space of almost complex structures that are compatible with* (E_π, ω)*. Then* $\mathcal{J}(E_\pi, \omega)$ *is non-empty and contractible.*

Before proving this theorem, we will make note that a pair of compatible structures (ω, J) on E_π induce an orthogonal structure on E_π. Define the bilinear structure g_J on E_π by $g_J(v, w) = \omega(v, Jw)$. For any $v, w \in \pi^{-1}(x)$ we have that

$$g_J(v, w) = \omega(v, Jw) = \omega(Jv, -w) = -\omega(-w, Jv) = \omega(w, Jv) = g_J(w, v)$$

and so g_J is symmetric. Additionally for any nonzero $v \in \pi^{-1}(x)$ we have that

$$g_J(v, v) = \omega(v, Jv) > 0$$

and thus g_J is non-degenerate and positive definite. Thus, g_J is an orthogonal structure.

Proof. Let $\mathfrak{Met}(E_\pi)$ denote the set of all orthogonal structures that can be placed upon E_π. It was proved in the proof of Proposition 11 that $\mathfrak{Met}(E_\pi)$ is nonempty. Suppose that there exists a continuous map $r : \mathfrak{Met}(E_\pi) \to \mathcal{J}(E_\pi, \omega)$ such that $r(g_J) = J$ for all $J \in \mathcal{J}(E_\pi, \omega)$ and $r(\Phi^* g) = \Phi^* r(g)$ for all symplectomorphisms Φ and $g \in \mathfrak{Met}(E_\pi)$.

Then, in particular, $\mathcal{J}(E_\pi, \omega)$ is non-empty. Moreover, if we fix $J_0 \in \mathcal{J}(E_\pi, \omega)$, the map $f_t : \mathcal{J}(E_\pi, \omega) \times I \to \mathcal{J}(E_\pi, \omega)$ defined by

$$f_t(J) = r((1 - t)g_J + t g_{J_0})$$

defines a homotopy from the identity map to the constant map.

We have therefore reduced the problem to constructing the retraction r. Fix $g \in \mathfrak{Met}(E_\pi)$. Since both g and ω are non-degenerate, there exists an automorphism $A : E_\pi \to E_\pi$ satisfying $\omega(v, w) = g(Av, w)$. Let A^* denote the skew-adjoint of A with respect to g. Since

$$g(Av, w) = \omega(v, w) = -\omega(w, v) = -g(Aw, v) = -g(v, Aw),$$

we have that $A^* = -A$. Thus $A^*A = -A^2$ is positive definite and self-adjoint with respect to g. Denote A^*A by P.

Over each contractible open subset of M, we can choose $2n$ linear independent sections of E_π, and we can then express $P = M^{-1}DM$, where the matrix of D with respect to the basis provided by these sections is diagonal. Since P is positive definite, the entries of the matrix of D over each fibre are positive. We can thus define $D^{\frac{1}{2}}$ to have the positive square root of the values of D along its diagonal. Thus, $Q = M^{-1}D^{\frac{1}{2}}M$ is a positive definite automorphism of E_π that is self-adjoint with respect to g.

Define the automorphism $J_g = Q^{-1}A$.

For each $x \in M$, the decomposition of $\pi^{-1}(x)$ into the eigenspaces of A, P and Q are identical. Thus, these automorphisms pairwise commute. From this,

$$J_g^2 = Q^{-1}AQ^{-1}A = -Q^{-2}A^*A = -I$$

and so J_g is a complex structure.

Moreover,

$$\omega(J_gv, J_gw) = \omega(Q^{-1}Av, Q^{-1}Aw) = \omega(v, Q^{-2}A^*Aw) = \omega(v, w)$$

and

$$\omega(v, J_gv) = g(Av, Q^{-1}Av) = g(Pv, Q^{-1}v) > 0$$

since both P and Q are positive definite.

Thus $J_g \in \mathcal{J}(E_\pi, \omega)$. Define $r : \mathfrak{Met}(E_\pi) \to \mathcal{J}(E_\pi, \omega)$ by $r(g) = J_g$. For $J \in \mathcal{J}(E_\pi, \omega)$, we have that

$$\omega(v, w) = g_J(Av, \omega) = \omega(Av, Jw) = -\omega(JAv, w)$$

for all $v, w \in \pi^{-1}(x)$ and so $A = J$. Thus, $P = Q = id$ and so $r(g_J) = J$. Additionally, for $r(\Phi^*g)$, A changes to $\Phi^{-1}A\Phi$ and thus

$$\Phi^*r(g) = J_{\Phi^*g} = \Phi^{-1}J_g\Phi = r(\Phi^*g).$$

Thus, r is the desired map, which completes this proof. □

Theorem 56. [26] *Let J be a complex structure on E_π. There exists a symplectic structure on E_π that is compatible with J. Moreover, the space of such structures is contractible.*

Proof. Let $g \in \mathfrak{Met}(E_\pi)$. Define a bilinear form on form ω_g on E_π by $\omega_g(v, w) = g(v, -Jw)$. Note that

$$\omega_g(v, w) = g(v, -Jw) = g(-Jw, v) = -g(w, -Jv) = -\omega_g(w, v)$$

and $\omega_g(v, Jv) > 0$ for all $v \neq 0$, showing that ω_g is symplectic. It then follows that

$$\omega_g(Jv, Jw) = g(Jv, w) = g(v, -Jw) = \omega_g(v, w)$$

and

$$\omega_g(v, Jv) = g(v, v) > 0$$

and so J is compatible with ω_g, showing that the space of such symplectic forms is non-empty.

Now suppose that both ω_0 and ω_1 are symplectic structures compatible with J. For $0 \leq t \leq 1$, define $\omega_t = (1-t)\omega_0 + t\omega_1$. Fix $x \in M$. It is immediate that ω_t is skew-symmetric and satisfies both $\omega_t(Jv, Jw) = \omega_t(v, w)$ for all $v, w \in \pi^{-1}(x)$ and $\omega_t(v, Jv) > 0$ for all $v \in \pi^{-1}(x)\backslash\{0\}$. In particular, the fact that $\omega_t(v, Jv) > 0$ for all $v \in \pi^{-1}(x)$ guarantees that ω_t is non-degenerate and thus a symplectic structure on E_π. We therefore have that the space of symplectic forms on E_π that are compatible with J form a convex, and thus contractible, space. □

In particular, we have the following theorem.

Theorem 57. [26] *Two symplectic vector bundles are isomorphic if and only if their underlying almost complex vector bundles are isomorphic.*

To study symplectic vector bundles, it is therefore sufficient to study complex vector bundles. One benefit of studying symplectic vector bundles via the related complex vector bundles is that we may define the Chern classes, an invariant of complex vector bundles, for symplectic vector bundles. We will define the characteristic classes as follows.

Theorem 58. [26] *There exists a unique functor c_1 (the first Chern number) from the category of symplectic vector bundles over a Riemann surface to $\mathbb{Z}$ satisfying the following axioms.*

- *The functor is faithful.*
- *Given a smooth map $f : \Sigma \to \Sigma'$ and symplectic vector bundle $\pi : E \to \Sigma'$, $c_1(f^*E) = \deg f \cdot c_1(E)$.*
- *Given symplectic vector bundles $\pi_i : E_i \to \Sigma$, $c_1(E_1 \oplus E_2) = c_1(E_1) + c_1(E_2)$.*

- $c_1(T\Sigma) = \chi(\Sigma)$.

The proof of this theorem is found on page 73 of [26].

Consider a symplectic vector bundle $\pi : E \to M$, where M is not necessarily a Riemann surface. To every smooth map $f : \Sigma \to M$ from Riemann surface Σ to M, we can associate the integer $c_1(f^*E)$. Moreover, if two such maps f_1 and f_2 are homologous, their degrees equal, and thus $c_1(f_1^*E) = c_1(f_2^*E)$. Thus, the Chern numbers associate a cohomology class $c_1(E) \in H^2(M;\mathbb{Z})$ to each bundle.

Of particular interest is the Chern class for the tangent bundle of a manifold. We will refer to this as the *canonical class* of the manifold, and denote it as K_M.

We will later perform calculations involving the canonical class of $\Sigma_g \times \Sigma_h$. Let π_1 and π_2 denote the standard projections to Σ_g and Σ_h respectively. Since $T(\Sigma_g \times \Sigma_h) \cong T\Sigma_g \times T\Sigma_h$, we have that $T(\Sigma_g \times \Sigma_h) \cong \pi_1^*T\Sigma_g \oplus \pi_2^*T\Sigma_h$. We therefore have that

$$\begin{aligned}(K_{\Sigma_g\times\Sigma_h})_{PD} &= c_1(\pi_1^*T\Sigma_g)_{PD} + c_1(\pi_2^*T\Sigma_h)_{PD} \\ &= (2-2g)[\{pt\}\times\Sigma_h] + (2-2h)[\Sigma_g \times \{pt\}].\end{aligned}$$

The canonical class of a manifold can be used to determine the genus of symplectic submanifolds of a symplectic 4-dimensional manifold via the *adjunction formula*. To understand this formula, we must also understand the intersection form of a 4-dimensional manifold.

Definition 59. *Given a 4-manifold M, the **intersection form** is the $\mathbb{Z}$-bilinear map $Q_M : H^2(M;\mathbb{Z}) \times H^2(M;\mathbb{Z}) \to \mathbb{Z}$ defined as $Q_M(\alpha,\beta) = \langle \alpha \cup \beta, [M]\rangle$, where $[M]$ is the fundamental class of M.*

Given the Poincaré isomorphism $H^2(M;\mathbb{Z}) \to H_2(M;\mathbb{Z})$, we may sometimes confuse the intersection form as acting on $H^2(M;\mathbb{Z}) \times H_2(M;\mathbb{Z})$ or $H_2(M;\mathbb{Z}) \times H_2(M;\mathbb{Z})$. Since Q_M is bilinear, the annihilator of Q_M contains the torsion component of $H^2(M;\mathbb{Z})$. We may therefore consider Q_M to be acting on $H^2(M;\mathbb{Z})$ modulo the torsion subgroup. We therefore have that the intersection form extends to $H^2(M;\mathbb{R}) \supset H^2(M;\mathbb{Z})/\text{Torsion}$. Note that in DeRham cohomology, $Q_M(\alpha,\beta)$ can be reinterpreted as $\int_M \alpha \wedge \beta$.

We may provide another characterization of the intersection form that is more deserving of the name. Given $\alpha, \beta \in H^2(M)$, let $a, b \in H_2(M)$ be their respective Poincaré duals. Let Σ_a and Σ_b be surfaces in M that represent a

and b respectively (recall that Σ represents the homology class $i_*([\Sigma])$ where $[\Sigma]$ is the fundamental class of Σ and i is the provided embedding of Σ in M; we will also denote this homology class as $[\Sigma]$). Moreover, suppose that Σ_a and Σ_b only intersect transversely, and that $\Sigma_a \cap \Sigma_b$ is finite. For each point $p \in \Sigma_a \cap \Sigma_b$, we will associate ± 1 depending on whether the natural isomorphism $T_p\Sigma_a \oplus T_p\Sigma_b \to T_pM$ is orientation-preserving ($+1$) or reversing (-1). We then have the following theorem.

Theorem 60. [16] *Let α, β, a, b, Σ_a and Σ_b be as above. We then have that*

$$Q_M(\alpha, \beta) = \sum_{\Sigma_a \cap \Sigma_b} \pm 1.$$

Proof. Let η be a smooth 2-form that represents α under the natural map $H^2(M;\mathbb{Z}) \to H^2(M;\mathbb{R})$. Assume that η is supported on a (small) tubular neighbourhood of Σ_a (the definition and proof of existence of a tubular neighbourhood are provided later in Definition 63 and Theorem 64 respectively). Similarly, choose θ to be 2-form that represents β.

Given $p \in \Sigma_a \cap \Sigma_b$, choose a neighbourhood U_p of p that admits local coordinates (x, y, z, w) such that $\Sigma_a = \{x = y = 0\}$ and $\Sigma_b = \{z = w = 0\}$. We may then assume that within this neighbourhood, η and θ are modelled as $\eta = f(x, y)dx \wedge dy$ and $\theta = f(z, w)dz \wedge dw$ where $f : \mathbb{R}^2 \to \mathbb{R}$ is a bump function centred around 0 with $\int_{\mathbb{R}^2} f = 1$. We then have that

$$Q_M(\alpha, \beta) = \int_M \eta \wedge \theta = \sum_{p \in \Sigma_a \cap \Sigma_b} \int_{U_p} \eta \wedge \theta.$$

The integral $\int_{U_p} \eta \wedge \theta$ can be rewritten as

$$\int_{\mathbb{R}^4} f(x, y)f(z, w)dx \wedge dy \wedge dz \wedge dw,$$

which is ± 1 depending on whether the basis $(\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z}, \frac{\partial}{\partial w})$ of T_pM is positively or negatively oriented. This completes the proof. □

Every 2-dimensional symplectic submanifold of a compact symplectic 4-dimensional manifold M satisfies an equation that provides the genus of the submanifold in terms of its self-intersection and its intersection with the canonical form. Thus equation is referred to as the *adjunction formula*.

Theorem 61. [26] *Let M be a compact 4-dimensional symplectic manifold. Let $i : X \to M$ be a symplectic submanifold. We then have the following equality:*

$$\chi(X) = Q_M(K_M, [X]) - Q_M([X], [X]).$$

A proof of this theorem is found on page 281 of [31].

Lastly, we will consider the existence of branched covers over closed oriented 4-manifolds.

Theorem 62. [20] *Let M be a closed oriented 4-manifold. Let X be a surface in M. Let $\xi \in H^2(M; \mathbb{Z})$ be Poincaré dual to $[X]$. Suppose furthermore that ξ is divisible by n (there exists $\alpha \in H^2(M; \mathbb{Z})$ such that $\xi = n \cdot \alpha$). There exists a cyclic n-fold branched cover of M with branch locus X.*

Proof. Since the Chern map $\check{H}^1(M, U_1) \to H^2(M)$ is an isomorphism, there exists a line bundle L_π with first Chern class ξ. We can then find a non-zero section $\sigma : M \to L_\pi$ that vanishes precisely on X. Since $\xi = n\alpha$, the line bundle K_ρ with first Chern class α satisfies $K_\rho^n = L_\pi$. Define $f : K_\rho \to L_\pi$ by $f(v) = v \otimes \cdots \otimes v$. Let $N = f^{-1}\sigma(M)$.

Given $x \in M \setminus X$, $\sigma(x)$ is a non-zero element of $\pi^{-1}(x)$, and so $\sigma(x)$ has n distinct roots that admit a natural $\mathbb{Z}_n$-torsor structure. This therefore extends to a natural $\mathbb{Z}_n$ action on $N \setminus f^{-1}\sigma(X)$ whose quotient is $M \setminus X$. Moreover, for any $x \in X$, $\sigma(x) = 0$ and thus this action degenerates on X. We thus have that $\pi(f|_N) : N \to M$ is the desired branched cover. $\square$

4.5 Gompf's Symplectic Sum

We begin this section by studying the restriction of symplectic structures to open subsets about submanifolds.

Definition 63. *Let M be a manifold with closed submanifold X. A* ***tubular neighbourhood*** *of X is an open neighbourhood U of X together with a diffeomorphism $\exp^{-1} : U \to \mathcal{N} \subseteq NX$, where $\mathcal{N}$ is a neighbourhood of the zero section of NX such that $\mathcal{N} \cap \pi_{NX}^{-1}(x)$ is convex for each $x \in X$.*

Note that for a tubular neighbourhood $(U, \exp^{-1})$ of $X \subset M$, $\exp^{-1}$ gives U a fibre bundle structure with base X and fibre the open ball of dimension $\dim M - \dim X$.

We will begin by proving the existence of a tubular neighbourhood for any compact X in M. This is commonly referred to in literature as the *Tubular Neighbourhood Theorem*

Theorem 64. [9] *Suppose that X is a k-dimensional compact submanifold of an n-dimensional manifold M. There exists a tubular neighbourhood of X.*

Proof. Fix a Riemannian metric g on M. For any $v \in T_pM$ there exists a unique geodesic whose derivative at p is v (cf. Theorem 5.8 of Chapter VII in [6]). Identifying tangent vectors with their corresponding geodesics, we can define $\exp : T_pM \to M$ by sending v to $v(1)$. Given any $p \in M$, there exists an open neighbourhood U of p and an $\epsilon > 0$ such that exp is defined and continuous on $\{X_q \in T_qM \mid q \in U, \|X_q\| < \epsilon\}$ (cf. [6]). Thus, since X is compact, we can find an open neighbourhood $\mathcal{U} \subset TM$ of the zero section of X such that exp is defined on $\mathcal{U}$. Consider the restriction of exp to $\mathcal{U} \cap NX$. For any $x \in X$, $T_{(x,0)}(\exp) : T_{(x,0)}NX \to T_xM$ is invertible, and so there exists an open neighbourhood about x such that exp is invertible. We can therefore define $\exp^{-1}$ on an open subset about X. We may assume that this neighbourhood retracts to X, and so this subset, along with $\exp^{-1}$, forms a tubular neighbourhood about X. □

Given two symplectic structures on the ambient manifold, we can now demonstrate the existence of symplectomorphisms between tubular neighbourhoods of a submanifold when the criterion below is met. The following is referred to as *Moser's Theorem*.

Theorem 65. [9] *Let M be a manifold that admits a submanifold X. Let ω_0 and ω_1 be symplectic forms on M such that $\omega_0|_p = \omega_1|_p$ for all $p \in X$ (as bilinear forms on T_pM). Then there exist open neighbourhoods U_0 and U_1 of X and a diffeomorphism $\psi : U_0 \to U_1$ such that $\psi(p) = p$ for all $p \in X$ and $\psi^*\omega_1 = \omega_0$.*

Proof. Let U_0 be an open tubular neighbourhood of X. Let $i : X \to U_0$ be the inclusion, and let $\pi : U_0 \to X$ be the bundle map induced from the normal bundle of X. Suppose that there exists a path ω_t, $t \in [0, 1]$, of symplectic 2-forms on U_0 such that there exists a 1-form μ satisfying $\frac{d}{dt}\omega_t = d\mu$ for all t. Suppose further that there exists a time-dependent vector field x_t on U_0 such that $\mu + \iota_{x_t}\omega_t = 0$ and $x_t(p) = 0$ for all $p \in X$.

Using Picard's Existence Theorem, for each $p \in X$ there exists an open neighbourhood of p and a unique isotopy ψ_t on this neighbourhood such that

$\frac{d}{dt}\psi_t = x_t \circ \psi_t$ and $\psi_0 = id$. Since X is compact, we can extend these isotopies to an isotopy defined over a tubular neighbourhood of X contained in U_0. Relabel this new tubular neighbourhood as U_0. Note that since $x_t(p) = 0$ for all $p \in X$, we have that $\psi_t(p) = p$ for all t.

We then have that

$$\begin{aligned}
0 &= \psi_t^*(d(\mu + \iota_{x_t}\omega_t)) \\
&= \psi_t^*(\frac{d}{dt}\omega_t + d\iota_{x_t}\omega_t) \\
&= \psi_t^*(\frac{d}{dt}\omega_t + d\iota_{x_t}\omega_t + \iota_{x_t}d\omega_t) && \text{(since } \omega_t \text{ is closed)} \\
&= \psi_t^*(\frac{d}{dt}\omega_t + \mathcal{L}_{x_t}\omega_t) && \text{(Cartan's magic formula)} \\
&= \frac{d}{dt}(\psi_t^*\omega_t). && \text{(cf. Theorem 6.4 of [9])}
\end{aligned}$$

In particular, $\psi_t^*\omega_t = \omega_0$ for all t. We can then set ψ to equal ψ_1 and U_1 to be the domain of ψ in order to fulfill the conclusion of the theorem.

It therefore suffices to find the above mentioned path ω_t, the 1-form μ and the time-dependent vector field x_t.

The two-form $\omega_1 - \omega_0$ is closed on U_0 and satisfies $(\omega_1 - \omega_0)|_p = 0$ for all $p \in X$. Moreover, $\omega_1 - \omega_0$ is exact. Indeed, note that $i^* : H^2(U_0; \mathbb{R}) \to H^2(X; \mathbb{R})$ is a group isomorphism (since i is a homotopy equivalence with homotopy inverse π) and that $[i^*(\omega_1 - \omega_0)] = 0$ since i^* is the restriction map. However, we require the stronger statement that there exists a $\mu \in \Omega^1(U_0)$ such that $d\mu = \omega_1 - \omega_0$ and that $\mu_p = 0$ for all $p \in X$. To prove this, we will define a chain homotopy between the chain maps id_{U_0} and $(i \circ \pi)^*$. Recall that a chain homotopy between id_{U_0} and $(i \circ \pi)^*$ is a sequence of group morphisms $Q : \Omega^i(U_0) \to \Omega^{i-1}(U_0)$ that satisfy

$$Qd + dQ = (i \circ \pi)^* - id$$

for all i. In the following diagram, this definition can be reinterpreted as stating that the difference of paired vertical maps is given by the sum along the two paths of the neighbouring parallelogram.

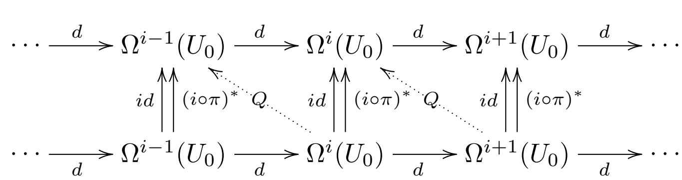

Define a retract ρ of U_0 onto X by defining $\rho_t : U_0 \to U_0$ by $\rho_t(p, v) = (p, tv)$ for $0 \le t \le 1$. This is well-defined by identifying U_0 with a convex neighbourhood of the zero-section of the normal bundle of X. For a given k-form, we can define $Q\omega = \int_0^1 \rho_t^*(\iota_{v_t}\omega)dt$ where $v_t(\rho_t(p, v))$ is tangent to the curve $\rho_s(p, v)$ at $s = t$. To see that Q is a chain homotopy, note that

$$\begin{aligned}
Qd\omega + dQ\omega &= \int_0^1 \rho_t^*(\iota_{v_t}d\omega)dt + d\int_0^1 \rho_t^*(\iota_{v_t}\omega)dt \\
&= \int_0^1 \rho_t^*(\iota_{v_t}d\omega + d\iota_{v_t}\omega)dt \\
&= \int_0^1 \rho_t^*(\mathcal{L}_{v_t}\omega)dt && \text{(Cartan's magic formula)} \\
&= \int_0^1 \frac{d}{dt}\rho_t^*\omega dt && \text{(cf. Theorem 6.4 of [9])} \\
&= \rho_1^*(\omega) - \rho_0^*(\omega) \\
&= id(\omega) - (i \circ \pi)^*(\omega).
\end{aligned}$$

Now that we have a chain homotopy, note that

$$\begin{aligned}
\omega_1 - \omega_0 &= id(\omega_1 - \omega_0) - (i \circ \pi)^*(\omega_1 - \omega_0) \\
&= Qd(\omega_1 - \omega_0) + dQ(\omega_1 - \omega_0) \\
&= dQ(\omega_1 - \omega_0),
\end{aligned}$$

and so $Q(\omega_1 - \omega_0)$ is a particular 1-form whose differential is $\omega_1 - \omega_0$. Moreover, since $\rho_t(p) = p$ for all $p \in X$, $v_t|_p = 0$, and so $Q(\omega_1 - \omega_0)$ is identically zero on X, showing that $Q(\omega_1 - \omega_0)$ is a suitable choice for μ.

Consider the line segment $\omega_t = (1-t)\omega_0 + t\omega_1 = \omega_0 + td\mu$, sitting inside the space of 2-forms. Since $\det \circ\, \omega_t : U_0 \times I \to \mathbb{R}$ is non-zero on the compact subspace $X \times I$, there exists an open neighbourhood of $X \times I$ inside $U_0 \times I$ such that $\det \circ\, \omega_t$ is non-zero. Therefore, we may shrink U_0 so that each ω_t is non-degenerate, and hence symplectic, on U_0.

Now, since the ω_t are non-degenerate, we can locally solve $\iota_{x_t}\omega_t = -\mu$ for a unique x_t. Since, ω_t is a smooth, time-dependent form, x_t is a time-dependent vector field. Note that $x_t(p) = 0$ whenever $\mu_p = 0$, and thus $x_t(p) = 0$ for all $p \in X$.

We therefore have suitable choices for ω_t, x_t and μ, completing the proof. □

Using Moser's Theorem, given two symplectic manifolds we can construct symplectomorphisms between tubular neighbourhoods of symplectic submanifolds when certain criteria are met. The corollary below is referred to in literature as the *Symplectic Neighbourhood Theorem*.

Corollary 66. [9] *For $i = 0, 1$, let (M_i, ω_i) denote a $2n$-dimensional symplectic manifold with compact symplectic submanifold $X_i \subset M_i$. Suppose that there exists a symplectomorphism $\varphi : X_0 \to X_1$ that can be extended to a symplectic vector bundle isomorphism $\Phi : NX_0 \to NX_1$ (regarding X_i as the zero section of the corresponding vector bundle). Fix such an extension of φ. Then φ extends to a symplectomorphism between tubular neighbourhoods $\widehat{\varphi} : \mathcal{N}(X_0) \to \mathcal{N}(X_1)$ with the property that $d\widehat{\varphi} : T_pM_0 \to T_{\widehat{\varphi}(p)}M_1$ restricts to Φ on $NX_0 = T_pX_0^{\omega}$.*

Proof. Using the Tubular Neighbourhood Theorem, we can extend φ to a diffeomorphism $\varphi' = \exp_1 \circ \Phi \circ \exp_0^{-1} : \mathcal{N}(X_0) \to \mathcal{N}(X_1)$ with the desired property. Since Φ is a symplectic vector bundle isomorphism, the symplectic forms ω_0 and $(\varphi')^*\omega_1$ on $\mathcal{N}(X_0)$ agree on X. By Theorem 65 there exist tubular neighbourhoods $U_0, U_1 \subset \mathcal{N}(X_0)$ of X_0 and a symplectomorphism $\psi : U_0 \to U_1$ such that $\psi^*((\varphi')^*\omega_1) = \omega_0$. We therefore have that $\varphi' \circ \psi : U_0 \to \varphi'(U_1)$ is the desired symplectomorphism. □

We now have the machinery necessary to perform the symplectic sum operation.

Let (M_1, ω_1) and (M_2, ω_2) be symplectic $2n$-dimensional manifolds. Suppose that there exists a symplectic $(2n-2)$-dimensional manifold X that is realized as an embedded symplectic submanifold of both M_1 and M_2 via the embeddings $\varphi_i : X \to M_i$. Moreover, suppose that the normal bundles of these two submanifolds are both trivial (implying that the realizations of X as a submanifold have self-intersection 0). From Corollary 66, there exist symplectomorphisms f_i from tubular neighbourhoods $X_i \times B^2_\epsilon$ of the zero sections of NX_i into tubular neighbourhoods of X_i in M_i. Moreover, the f_i restrict to the canonical maps between the zero sections and X_i.

Fix $\delta < \epsilon$. We can find an orientation-preserving symplectomorphism ψ from the annulus $A_{\epsilon,\delta} = B^2_\epsilon \setminus B^2_\delta$ onto itself, which maps the ϵ boundary component onto the δ boundary component, and vice-versa.

Note that it is at this step that we require the codimension of X (and hence the dimension of the fibre of the normal bundle) to be 2: if such a symplectomorphism $B^{2n}_\epsilon \setminus B^{2n}_\delta \to B^{2n}_\epsilon \setminus B^{2n}_\delta$ exists, we can glue two copies of B^{2n}_ϵ together along this map to achieve a symplectic structure on S^{2n}. Since the only sphere that admits a symplectic structure is S^2 (since the second cohomology of a compact symplectic manifold is non-zero), n must be 1.

We can now define $M_1 \#_X M_2 = (M_1 \setminus U_1) \cup_\varphi (M_2 \setminus U_2)$, identifying the two components using $\varphi = (\varphi_2 \circ \varphi_1^{-1}) \times \psi$. The symplectic forms ω_1 and ω_2 agree on $X \times A_{\epsilon,\delta}$, and thus define a 2-form ω on $M_1 \#_X M_2$. Moreover, since every point admits a neighbourhood such that ω agrees with either ω_1 or ω_2, this 2-form is symplectic. Note that $M_1 \#_X M_2$ depends on the choice of ψ.

We have therefore proven the following theorem.

Theorem 67. [26] *For $i = 1, 2$, let (M_i, ω_i) be a $2n$-dimensional manifold that admits a compact $(2n-2)$-dimensional symplectic submanifold X_i that has trivial normal bundle. Suppose furthermore that (X_0, ω_0) and (X_1, ω_1) are symplectomorphic. Then, the above described manifold $M_1 \#_X M_2$ admits a symplectic structure.*

In the ensuing application of symplectic summing, we will be considering submanifolds N_i of M_i that transversely meet the submanifolds X_i in the same number of points. We then require a method to line up these points in such a way that we obtain a submanifold of $M_1 \#_X M_2$. The next lemma shows that we can find a symplectomorphism that is isotopic to the given map between X_0 and X_1 that lines up the points of intersection with the N_i manifolds, thus allowing for the N_i to glue together when taking the symplectic sum of the M_i's. The corollary after provides sufficient criteria for the submanifold $N_0 \#_X N_1 \subset M_0 \#_X M_1$ to be symplectic.

Lemma 68. [14] *Let (Σ, ω) and (Σ', ω') be closed, connected, symplectic surfaces with the same area. Let $x_1, \ldots, x_n \in \Sigma$, and let $y_1, \ldots, y_n \in \Sigma'$. Every area-preserving diffeomorphism $f : \Sigma \to \Sigma'$ is isotopic to a symplectomorphism that maps x_i to y_i.*

Proof. Note that we may assume that $f(x_i) = y_i$ for each i, since we can smoothly isotope f to satisfy this requirement. Let ω_0 and ω_1 denote the symplectic forms ω and $f^*\omega'$ of Σ respectively. As in Theorem 65, we are in a position to apply Moser's argument. We wish to find an isotopy between f and a symplectomorphism that is constant on the x_i. Suppose we had a path of symplectic forms ω_t from ω_0 to ω_1. Suppose further that the time-derivative $\frac{d}{dt}\omega_t$ was an exact form $d\sigma$. As in the proof of Theorem 65 we could then solve

$\iota_{v_t}\omega_t = -\sigma$ for the time-dependent vector field v_t, and solve $\frac{d}{dt}\psi_t = v_t \circ \psi_t$ for an isotopy ψ_t satisfying $\psi_t^*\omega_t = \omega_0$ and $\psi_0 = id$. We would then have that f is isotopic to $f \circ \psi_1$ where $(f \circ \psi_1)^*\omega' = \psi_1^*\omega_1 = \omega_0$, showing that $f \circ \psi_1$ is a symplectomorphism. By suitably choosing σ we can guarantee that this isotopy is constant on the x_i's.

Define $\eta = \omega_1 - \omega_0$. Since f is area-preserving, $f^* : H^2(\Sigma'; \mathbb{R}) \to H^2(\Sigma; \mathbb{R})$ maps the preferred generator of $H^2(\Sigma'; \mathbb{R})$ to the preferred generator of $H^2(\Sigma; \mathbb{R})$. We therefore have that $[\eta] = 0$, and thus there exists $\varphi \in \Omega^1(\Sigma)$ such that $d\varphi = \eta$. Choose a smooth $g : \Sigma \to \mathbb{R}$ such that $dg|_{x_i} = \varphi|_{x_i}$ for all $1 \leq i \leq n$. We then have that $\varphi' = \varphi - dg$ satisfies $d\varphi' = \eta$ and $\varphi'|_{x_i} = 0$ for all i.

Define $\omega_t = \omega_0 + t\eta$ for $t \in [0, 1]$. Since f is area preserving, we have that the signs of $\omega_0(x, y)$ and $\omega_1(x, y)$ agree for any oriented basis $\{x, y\}$ of $T_p\Sigma$. Thus, for such a basis, $\omega_t(x, y) = (1 - t)\omega_0(x, y) + t\omega_1(x, y)$ is non-zero. In particular, ω_t is non-degenerate everywhere, and hence symplectic.

Moreover, we have that $\frac{d}{dt}\omega_t = \eta = d\varphi'$, which is exact. As mentioned above, we can now solve for time-dependent vector field v_t and isotopy ψ_t. Note that since ω_t is everywhere non-degenerate, and $\iota_{v_t}\omega_t|_{x_i} = -\varphi'|_{x_i} = 0$, $v_t|_{x_i} = 0$ for all t. We therefore have that $\frac{d}{dt}\psi_t|_{x_i} = 0$ and thus $\psi_t(x_i) = \psi_0(x_i)$ for all t and $1 \leq i \leq n$. In particular, $f \circ \psi_1$ is the desired symplectomorphism. □

This lemma implies the following corollary.

Corollary 69. [14] *For $i = 0, 1$, let (M_i, ω_i) be a 4-dimensional symplectic manifold, and let X_i be as in Theorem 67. Suppose that for each i there exists a symplectic surface F_i of M_i that intersects X_i transversely in a finite number of points. Suppose furthermore that $|F_0 \cap X_0| = |F_1 \cap X_1|$. Then, given a bijection $f : F_0 \cap X_0 \to F_1 \cap X_1$ there exists a symplectic surface $F \subset M_0 \#_X M_1$ that restricts to F_i in $M_i \setminus U_i$ and connects these two components according to f.*

Chapter 5

Generalization of Smith's construction

5.1 The Construction

Fix $h \geq 0$. Let $\{x_1, \ldots, x_p\}$ denote the basepoint used in the construction of $\mathrm{Br}_p(\Sigma_h)$. Recall from Section 3.2 that a braid $\beta \in \mathrm{Br}_p(\Sigma_h)$ admits a decomposition into p strands $\beta_i : I \to \Sigma_h$ such that for all $1 \leq i \leq p$, both $\beta_i(0)$ and $\beta_i(1)$ are in $\{x_1, \ldots, x_p\}$. Considering the group morphism $\sigma : \mathrm{Br}_p(\Sigma_h) \to S_p$ described in Theorem 26, we can identify the ending point of the graph of β_i with the starting point of $\beta_{\sigma(\beta)(i)}$, thus defining a unique isotopy class of embeddings of a set of circles into $S^1 \times \Sigma_h$. Moreover, this set of circles is naturally indexed by the cycles of the disjoint cycle decomposition of $\sigma(\beta)$. Let us denote this submanifold by Γ_β. We can then define the embedded submanifold $T_\beta = S^1 \times \Gamma_\beta \subset T^2 \times \Sigma_h$. Note that T_β is a union of tori whose path components are again indexed by the cycles of $\sigma(\beta)$.

Since we have a preferred decomposition of T^2 as $S^1 \times S^1$, we have a preferred area form on T^2 that is locally defined by $d\theta \wedge d\varphi$, where θ and φ are local coordinates of the circles. Let η be any area form of Σ_h. The product symplectic form on $T^2 \times \Sigma_h$ is then locally described as $d\theta \wedge d\varphi + \omega_{\Sigma_h}$ of $T^2 \times \Sigma_h$. If T_β passes through the origin of these local coordinates, then $T_0(T_\beta)$ is spanned by

$\frac{\partial}{\partial\theta}$ and $\frac{\partial}{\partial\varphi} + v$ for some nonzero tangent vector v of $T_0(\Sigma_h)$. Thus,

$$\begin{aligned}\omega\left(\frac{\partial}{\partial\theta}, \frac{\partial}{\partial\varphi} + v\right) &= d\theta \wedge d\varphi\left(\frac{\partial}{\partial\theta}, \frac{\partial}{\partial\varphi} + v\right) + \omega_{\Sigma_h}\left(\frac{\partial}{\partial\theta}, \frac{\partial}{\partial\varphi} + v\right) \\ &= d\theta \wedge d\varphi\left(\frac{\partial}{\partial\theta}, \frac{\partial}{\partial\varphi}\right) = 1.\end{aligned}$$

Since the origin of these local coordinates was arbitrary, T_β is a symplectic submanifold.

Recall that a connected m-dimensional submanifold $i : X \to M$ is said to represent the homology class $i_*([X]) \in H_m(M)$, where $[X]$ is the fundamental class of X. We will also denote the homology class $i_*([X])$ as $[X]$. If X is not connected, say $X = \sqcup_{i=1}^n X_i$ where each X_i is connected, we will say that X represents the homology class $[X] = \sum_{i=1}^n [X_i]$. One can then compute $[T_\beta]$ in terms of β. We may assume that T_β is connected, as otherwise each component of T_β can be created using a braid with fewer strands than β. Consider the maps induced by the projections

$$pr_1 : \pi_1(S^1 \times \Sigma_h) \to \pi_1(S^1) \quad \text{and} \quad pr_2 : \pi_1(S^1 \times \Sigma_h) \to \pi_1(\Sigma_h).$$

Changing T_β via homotopy does not affect the represented homology class. We will define a homotopy of T_β by homotoping Γ_β into a curve $\lambda\#pr_2(\beta)$, such that $pr_1(\lambda) = p$ and $pr_2(\lambda) = id$. Define T_λ as $S^1 \times \mathrm{Im}(\lambda)$ and $T_{pr_2(\beta)}$ as $S^1 \times \mathrm{Im}(pr_2(\beta))$. It then follows that

$$[T_\beta] = [T_\lambda] + [T_{pr_2(\beta)}] = p[T^2 \times \{pt\}] + [S^1 \times \mathrm{Im}(pr_2(\beta))] \in H_2(T^2 \times \Sigma_h; \mathbb{Z}).$$

The submanifolds T_β can be used to construct symplectic submanifolds of $\Sigma_g \times \Sigma_h$ for $g \geq 1$. Fix two braids $\beta_1, \beta_2 \in \mathrm{Br}_p(\Sigma_h)$. Regard $T^2 \times \Sigma_h$ as the trivial fibre bundle over T^2 with fibre Σ_h. Any fibre of this bundle admits a trivial normal bundle (since we can describe tubular neighbourhoods about the fibre by $D^2 \times \Sigma_h$). We can therefore apply Gompf's symplectic summing process (Theorem 67) to glue two copies of this bundle together to obtain the symplectic manifold $\Sigma_2 \times \Sigma_h$. Moreover, since both T_{β_1} and T_{β_2} intersect each fibre in exactly p points, we can apply Corollary 69 to glue these two submanifolds together to obtain a symplectic submanifold $\Theta(\beta_1, \beta_2)$ of $\Sigma_2 \times \Sigma_h$. More generally, we may repeat this process any finite number of times to construct a symplectic submanifold $\Theta(\beta_1, \ldots, \beta_g) \subset \Sigma_g \times \Sigma_h$.

Using the above descriptions of $[T_\beta]$, we can describe $[\Theta(\beta_1, \ldots, \beta_g)]$. Let γ_i be a circle in Σ_g that was the first S^1 factor of the i^{th} copy of $T^2 \times \Sigma_h$ before performing Gompf's symplectic summing process. Since the surfaces $T^2 \times \{pt\}$ glue together under the symplectic sum to become $\Sigma_g \times \{pt\}$, we have that

$$[\Theta(\beta_1, \ldots, \beta_g)] = p[\Sigma_g \times \{pt\}] + \sum_{i=1}^{g} [\gamma_i \times \mathrm{Im}(pr_2 \circ \beta_i)] \tag{5.1}$$

in $H_2(\Sigma_g \times \Sigma_h; \mathbb{Z})$.

From this, it follows that $\Theta(\beta_1, \ldots, \beta_g)$ has self-intersection 0. To prove this, it suffices to show that

- $[\Sigma_g \times \{pt\}]^2 = 0$,
- $[\Sigma_g \times \{pt\}] \cdot [\gamma_i \times \mathrm{Im}(pr_2(\beta_i))] = 0$ for all $1 \leq i \leq g$,
- $[\gamma_i \times \mathrm{Im}(pr_2(\beta_i))] \cdot [\gamma_j \times \mathrm{Im}(pr_2(\beta_j))] = 0$ for all $1 \leq i, j \leq g$.

The first requirement is clear since any choice of two distinct points of Σ_h provides submanifolds representing $[\Sigma_g \times \{pt\}]$ that are disjoint. The second requirement can be satisfied by choosing a point of Σ_h that does not lie in $\mathrm{Im}(pr_2(\beta_i))$. The third requirement is trivial when $i \neq j$ since γ_i and γ_j are disjoint. When $i = j$ we can perturb γ_i to a non-intersecting copy contained within Σ_g, thus fulfilling the third and last requirement.

Since $\Theta(\beta_1, \ldots, \beta_g)$ has trivial self-intersection, the adjunction formula states that

$$\begin{aligned}
\chi(\Theta(\beta_1, \ldots, \beta_g)) &= K_{\Sigma_g \times \Sigma_h} \cdot [\Theta(\beta_1, \ldots, \beta_g)] \\
&= \Big((2-2h)[\Sigma_g \times \{pt\}] + (2-2g)[\{pt\} \times \Sigma_h]\Big) \cdot \\
&\qquad \Big(p[\Sigma_g \times \{pt\}] + \sum_{i=1}^{g} [\gamma_i \times \mathrm{Im}(pr_2(\beta_i))]\Big).
\end{aligned}$$

Using the bilinearity of the self-intersection form and the above-computed intersection numbers, we then have that

$$\begin{aligned}
\chi(\Theta(\beta_1, \ldots, \beta_g)) &= (2-2g)p([\Sigma_g \times \{pt\}] \cdot [\{pt\} \times \Sigma_h]) \\
&\qquad + \sum_{i=1}^{g} (2-2h)[\{pt\} \times \Sigma_h] \cdot [\gamma_i \times \mathrm{Im}(pr_2(\beta_i))].
\end{aligned}$$

The intersection between $[\Sigma_g \times \{pt\}]$ and $[\{pt\} \times \Sigma_h]$ is clearly 1, and $[\{pt\} \times \Sigma_h] \cdot [\gamma_i \times \mathrm{Im}(pr_2(\beta_i))]$ can be seen to be 0 by choosing the point not in γ_i. Thus $\chi(\Theta(\beta_1, \ldots, \beta_g)) = (2-2g)p$ and so $\Theta(\beta_1, \ldots, \beta_g)$ has genus equal to $k = p(g-1) + 1$.

5.2 Proof of Non-Isotopy

Suppose that there exist finite sequences of braids $(\alpha_1, \dots, \alpha_g)$ and $(\beta_1, \dots, \beta_g)$ such that Γ_{α_i} is homotopic to Γ_{β_i} for all i. These homotopies can then be used to construct homotopies between T_{α_i} and T_{β_i} for all i, which in turn can be used to construct a homotopy between $\Theta(\alpha_1, \dots, \alpha_g)$ and $\Theta(\beta_1, \dots, \beta_g)$. In particular, for a given $1 \leq j \leq g$, if we fix all β_i for $i \neq j$, then as we vary β_j among the local braids of $\mathrm{Br}_p(\Sigma_h)$, the resulting surfaces are pairwise homotopic. Furthermore, if we demand that β_j is always connected then these surfaces will also all be connected. We therefore have the following theorem.

Theorem 70. *Assume that either $h = 0$ and $p \geq 4$, or that $h \geq 1$ and $p \geq 2$. Let $g \geq 1$, and let $\gamma_1, \dots, \gamma_g$ be simple closed curves in Σ_g as in (5.1). Choose any $j \in \{1, \dots, g\}$ and let $I(j) = \{1, \dots, g\} \setminus \{j\}$. For each $i \in I(j)$, choose $\mu_i \in H_1(\Sigma_h; \mathbb{Z})$, and let $d_i = \mathrm{div}(\mu_i)$. Let n be the greatest common divisor of $\{p\} \cup \{d_i \mid i \in I(j)\}$. If $n \geq 2$, then the homology class*

$$p[\Sigma_g \times \{\mathrm{pt}\}] + \sum_{i \in I(j)} [\gamma_i] \times \mu_i \in H_2(\Sigma_g \times \Sigma_h; \mathbb{Z}) \tag{5.2}$$

contains an infinite family of homotopic but pairwise non-isotopic symplectic surfaces of genus equal to $p(g-1)+1$.

Proof. For each $i \in I(j)$, choose a curve β_i in $\mathrm{Br}_p(\Sigma_h)$ such that $p-1$ of the strands are trivial and the other strand, considered as an element in $\pi_1(\Sigma_h)$, maps to μ_i under the abelianization map $\pi_1(\Sigma_h) \to H_1(\Sigma_h; \mathbb{Z})$. Let $\Lambda_{p,h}$ denote the subset of $\mathrm{Br}_p(\Sigma_h)$ consisting of connected local braids, and choose arbitrary $\beta_j \in \Lambda_{p,h}$.

From the paragraphs preceding this theorem, it is clear that the submanifolds constructed as we vary β_j are pairwise homotopic.

Since n divides the homology class of $\Theta(\beta_1, \dots, \beta_g)$, Theorem 62 provides an n-fold branched cover $Z(\beta_1, \dots, \beta_g)$ of $\Sigma_g \times \Sigma_h$ with branch locus $\Theta(\beta_1, \dots, \beta_g)$ (up to diffeomorphism). Since an isotopy between any two Θ's determines a homeomorphism between the branched covers, it suffices to prove that the set $\{Z(\beta_1, \dots \beta_g) \mid \beta_j \in \Lambda_{p,h}\}$ is infinite.

In order to distinguish two such spaces, we will first realize these branched covers as genuine fibre bundles. Consider the morphism

$$\tilde{\pi} : Z(\beta_1, \dots, \beta_g) \xrightarrow{\pi} \Sigma_g \times \Sigma_h \xrightarrow{pr_1} \Sigma_g.$$

The fibre over a point in Σ_g is the preimage of $\{pt\} \times \Sigma_h$ under the projection of the branched cover $\pi : Z(\beta_1, \ldots, \beta_g) \to \Sigma_g \times \Sigma_h$. The restriction of π to this pre-image is clearly a cyclic branched cover over Σ_h with branch locus equal to the intersection of $\Theta(\beta_1, \ldots, \beta_g)$ with the fibre. This set has cardinality p, and so Corollary 22 implies that $\pi^{-1}(pt) \cong \Sigma_k$ for $k = hn + \frac{1}{2}(n-1)(p-2)$. Note that with the given setup, k can potentially be any positive integer.

Since Σ_g is both connected and aspherical, the long exact sequence associated to this bundle ends in

$$1 \longrightarrow \pi_1(\Sigma_k) \longrightarrow \pi_1(Z(\beta_1, \ldots, \beta_g)) \longrightarrow \pi_1(\Sigma_g) \longrightarrow 1.$$

We therefore have a natural group morphism $\pi_1(\Sigma_g) \to \mathrm{Out}(\pi_1(\Sigma_k))$ given as follows. $\pi_1(Z(\beta_1, \ldots, \beta_g))$ acts on $\pi_1(\Sigma_k)$ by conjugation. If we consider this action up to conjugation by $\pi_1(\Sigma_k)$ (the inner automorphisms) then $\pi_1(\Sigma_k)$ clearly acts trivially upon itself, and so this action descends to an action by

$$\pi_1(Z(\beta_1, \ldots, \beta_g))/\pi_1(\Sigma_k) \cong \pi_1(\Sigma_g)$$

(the monodromy representation).

Lemma 5.1 of [19] claims, in particular, that since Σ_k is aspherical and closed and Σ_g is closed, that $Z(\beta_1, \ldots, \beta_g)$ is determined by the conjugacy class of this action. It therefore suffices to construct an infinite number of non-conjugate group homomorphisms $\pi_1(\Sigma_g) \to \Gamma_0(\Sigma_k)$ that can be realized as such monodromies.

As an aside, note that if the base space M is a 2-dimensional finite CW-complex, then every group morphism $\varphi : \pi_1(M) \to \Gamma_0(F)$ is realizable as a monodromy representation. We may assume that the given CW-complex of M admits a single point as its 0-skeleton by contracting a maximal tree within the 1-skeleton. Then, 1-cells correspond to generators of $\pi_1(M)$ and 2-cells correspond to relations that we place upon those generators. Moreover, over each cell the fibre bundle is necessarily trivial. Thus, consider the trivial bundle over a 1-cell, and glue the two end fibres together via the image of the 1-cell under φ. Similarly glue the trivial bundles over the 2-cells along the bundle over the 1-skeleton according to the relation that the 2-cell corresponds to. The fact that $\pi_1(M)$ is a group morphism is equivalent to the fact that this bundle is well-defined.

Returning to the proof, we note that since $\tilde{\pi}$ restricts to the branched cover $\Sigma_k \to \Sigma_h$ on its fibres, the monodromy representation takes values in $\Gamma^{\zeta}(\Sigma_k)$

where ζ is the automorphism on Σ_k that induces this branched cover.

We will consider the quotient of the monodromy representation by the subgroup $\langle\zeta\rangle$. Recall from Theorem 48 that $\Gamma^{\zeta}(\Sigma_k)/\langle\zeta\rangle$ is isomorphic to $\Gamma_p(\Sigma_h)$.

Recall from Theorem 37 that there exists an injective group morphism $\Psi : \mathrm{Br}_p(\Sigma_h)/\Delta_{p,h} \to \Gamma_p(\Sigma_h)$, where $\Delta_{p,h}$ is the centre of $\mathrm{Br}_p(\Sigma_h)$. Considering the canonical presentation of $\pi_1(\Sigma_g)$ given by

$$\pi_1(\Sigma_g) = \langle \gamma_1, \dots \gamma_g, \delta_1, \dots, \delta_g \mid \prod [\gamma_i, \delta_i]_{i=1}^g \rangle$$

where the γ_i's are as in (5.1), we can describe the quotients of the monodromy homomorphism by

$$\gamma_i \mapsto 1, \quad \delta_i \mapsto \Psi(q_\Delta(\beta_i)) \qquad \text{for all } 1 \le i \le g$$

where q_Δ is the quotient of $\mathrm{Br}_p(\Sigma_h)$ by its centre.

It therefore suffices to show that the number of conjugacy classes of such group morphisms is infinite as we vary β_j. It therefore suffices to show that an infinite number of conjugacy classes of $\Gamma_p(\Sigma_h)$ non-trivially intersect $\{\Psi(q_\Delta(\beta)) \in \Gamma_p(\Sigma_h) \mid \beta \in \Lambda_{p,h}\}$.

First consider the case $h \ge 1$ and $p \ge 2$. Let D^2 be a disk that contains the marked points $\{x_1, \dots, x_p\}$. Using Corollary 4.2 of [29], we have that the embedding $D^2 \hookrightarrow \Sigma_h$ induces an injective map $\Gamma_p(D^2) \hookrightarrow \Gamma_p(\Sigma_h)$.

It is sufficient to find an infinite number of representatives of conjugacy classes of $\Gamma_p(D^2)$ that non-trivially intersect $\Psi(q_\Delta(\Lambda_{p,h}))$, for suppose that $[f_1], [f_2] \in \Gamma_p(D^2)$ and that there exists $[\varphi] \in \Gamma_p(\Sigma_h)$ such that $[\varphi^{-1}][f_1][\varphi] = [f_2]$. We can choose a representative φ of $[\varphi]$ that maps D^2 homeomorphically onto itself and restricts to the identity on ∂D^2. Define $\varphi' \in \Gamma_p(D^2)$ to equal φ on D^2 and equal the identity elsewhere. We then have that $\varphi'^{-1} f_1 \varphi' = f_2$, and so f_1 and f_2 are conjugate in $\Gamma_p(D^2)$. Thus, it suffices to find infinitely many conjugacy classes of $\Gamma_p(D^2)$ that non-trivially intersect $\Psi(q_\Delta(\Lambda_{p,h}))$.

It follows from (3.1) that, since $\pi_1(\mathrm{Aut}_0(D^2))$ and $\pi_0(\mathrm{Aut}_0(D^2))$ are trivial, $\mathrm{Br}_p(D^2) \cong \Gamma_p(D^2)$. Since two connected braids in D^2 represent the same conjugacy class only if they close to provide the same knot in S^3, it suffices to note that there exists an infinite family of knots in S^3 that can be constructed from p-strand braids (since $p \ge 2$).

The case when $h = 0$ and $p = 4$ is due to [25]. Recall that $\Gamma_4(S^2)$ is generated by the half-Dehn twists $\omega_1, \dots, \omega_3$, which are the image of the standard generators of the braid group. Consider the set

$$\{\omega_1^{2m+1}\omega_2\omega_3 \mid m \geq 0\}.$$

These mapping classes all induce the 4-cycle $(1,2,3,4)$ and thus are connected. Moreover, under the representation ρ described in Section 3.4, we have that

$$\rho(\omega_1^{2m+1}\omega_2\omega_3) = \pm \begin{bmatrix} 1 & 2m+1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} = \pm \begin{bmatrix} -2m & 1 \\ -1 & 0 \end{bmatrix}.$$

In particular, we have that $tr(\rho(\omega_1^{2m+1}\omega_2\omega_3)) = -2m$, demonstrating that as we vary m we realize different conjugacy classes of $\Gamma_4(S^2)$.

Lastly, we need to consider the case where $h = 0$ and $p \geq 5$. Again, we have that $\Gamma_p(S^2)$ is generated by the half-Dehn twists $\omega_1, \ldots, \omega_{p-1}$. Let ω_p denote the braid twist between points x_p and x_1. Consider the set

$$\{\omega_1^{2m+1}\omega_2 \cdots \omega_{p-1} \mid m \geq 0\}.$$

Note that these mapping classes all induce the same p-cycle $(1, \ldots, p)$ of S_p, and so are connected. Let $\xi = \omega_1 \cdots \omega_{p-1}$. If we consider the fixed points $x_1, \ldots, x_p$ to be equally distributed around the equator of the sphere, then ξ can be realized as the rotation of the sphere about the north-south axis by angle $\frac{2\pi}{p}$. The above set of mapping classes can be rewritten as

$$\{\omega_1^{2m}\xi \mid m \geq 0\}.$$

Using the above description of ξ we have that $\xi\omega_i^2\xi^{-1} = \omega_{i+1}^2$. This implies that $(\omega_1^{2m}\xi)^p = \omega_1^{2m}\omega_2^{2m}\cdots\omega_p^{2m}\xi^p = \omega_1^{2m}\cdots\omega_p^{2m}$. For $m_1, m_2 \in \mathbb{N}$, $\omega_1^{2m_1}\xi$ is in the same conjugacy class as $\omega_2^{2m_2}\xi$ only if $(\omega_1^{2m_1}\xi)^p$ and $(\omega_1^{2m_2}\xi)^p$ are in the same conjugacy class. It therefore suffices to show that these p^{th} powers represent infinitely many conjugacy classes.

Moreover, since these p^{th} powers lie within $\mathrm{P}\Gamma_p(S^2)$, and since $\mathrm{P}\Gamma_p(S^2)$ is a finite index subgroup of $\Gamma_p(S^2)$, it suffices to show that these elements represent infinitely many conjugacy classes of $\mathrm{P}\Gamma_p(S^2)$. For given a finite index subgroup H of G, suppose that there exists a subset $\{g_k\}_{k\in\mathbb{N}}$ of H that are pairwise non-conjugate in H but pairwise conjugate in G. Let $\{a_1, \ldots, a_n\}$ be right coset representatives of H in G. We then have that for any $k \in \mathbb{N}$ there exists $h \in H$ and $1 \leq i \leq n$ such that

$$g_1 = (ha_i)^{-1}g_k(ha_i) = a_i^{-1}(h^{-1}g_kh)a_i$$

and so

$$a_i g_1 a_i^{-1} = h^{-1} g_k h.$$

In particular, there must exist two $j, k \in \mathbb{N}$ such that g_j and g_k are both conjugate in H to the same $a_i g_1 a_i^{-1}$. However this implies that g_j and g_k are conjugate within H, a contradiction. It therefore suffices to show that the elements

$$\{\omega_1^{2m} \cdots \omega_p^{2m} \mid m \geq 0\}$$

are non-conjugate in $\mathrm{P}\Gamma_p(S^2)$.

For $p \geq 4$ there exists a group morphism $\mathrm{P}\Gamma_p(S^2) \to \mathrm{P}\Gamma_4(S^2)$ defined by simply 'forgetting' that the points $x_5, \ldots, x_p$ are fixed. Under this map ω_i maps to ω_i for $i = 1, 2, 3$ and ω_i maps to 1 for $i \geq 4$. We therefore have that the set of potential conjugacy class representatives maps to the set

$$\{\omega_1^{2m} \omega_2^{2m} \omega_3^{2m} \mid m \geq 0\}.$$

The above representation of $\Gamma_4(S^2)$ restricts to a representation of $\mathrm{P}\Gamma_4(S^2)$. Using that map we have that

$$\begin{aligned} \rho(\omega_1^{2m} \omega_2^{2m} \omega_3^{2m}) &= \pm \begin{bmatrix} 1 & 2m \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -2m & 1 \end{bmatrix} \begin{bmatrix} 1 & 2m \\ 0 & 1 \end{bmatrix} \\ &= \pm \begin{bmatrix} 1-4m^2 & 4m-8m^3 \\ -2m & 1-4m^2 \end{bmatrix}. \end{aligned}$$

We therefore have that $tr(\rho(\omega_1^{2m} \omega_2^{2m} \omega_3^{2m})) = 2 - 8m^2$ and so these elements are non-conjugate in $\mathrm{P}\Gamma_4(S^2)$, completing the proof. □

Bibliography

[1] E. Artin: Theorie der zöpfe, *Hamburg Univ. Math. Sem. Abhandl.* **4** (1926), 101–126.

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

[3] J. S. Birman: On braid groups, *Comm. Pure Appl. Math.* **22** (1969), 41–72.

[4] J. S. Birman: Mapping class groups and their relationship to braid groups, *Comm. Pure Appl. Math.* **22** (1969), 213–238.

[5] J. S. Birman: *Braids, Links, and Mapping Class Groups.* Annals of Mathematics Studies, No. 82, Princeton University Press, Princeton, NJ, 1974.

[6] W. M. Boothby: *An Introduction to Differentiable Manifolds and Riemannian Geometry.* Academic Press Inc., Montreal, 1986.

[7] R. Bott and L. W. Tu: *Differential Forms in Algebraic Topology.* Graduate Texts in Mathematics, No. 82, Springer-Verlag, New York-Heidelberg, 1982.

[8] K. S. Brown: *Cohomology of Groups.* Graduate Texts in Mathematics, No. 87, Springer-Verlag, New York-Heidelberg, 1982.

[9] A. Cannas da Silva: *Lectures on Symplectic Geometry.* Lecture Notes in Matematics, No. 1764, Springer-Verlag, New York-Heidelberg, 2001.

[10] A. L. Edmonds, R. S. Kulkarni and R. E. Stong: Realizability of branched coverings of surfaces, *Trans. Amer. Math. Soc.* **282** (1984), 773–790.

[11] E. Fadell and J. Van Buskirk: The braid groups of E^2 and S^2, *Duke Math. J.* **29** (1962), 243–257.

[12] R. Fintushel and R. J. Stern: Symplectic surfaces in a fixed homology class, *J. Differential Geom.* **52** (1999), 203–222.

[13] R. H. Fox and L. Neuwirth: The braid groups, *Math. Scand.* **10** (1962), 119–126.

[14] R. E. Gompf: A new construction of symplectic manifolds, *Ann. of Math.* **142** (1995), 527–595.

[15] P. Griffiths and J. Harris: *Principles of Algebraic Geometry.* John Wiley & Sons Inc., Toronto, 1978.

[16] R. E. Gompf and A. I. Stipsicz: 4-*Manifolds and Kirby Calculus.* Graduate Studies in Mathematics, Vol. 20, Amer. Math. Soc., Providence, RI, 1999.

[17] A. Hatcher: *Algebraic Topology.* Cambridge University Press, Cambridge, 2002.

[18] A. Hatcher: *Vector Bundles and K-theory.* preprint.

[19] J. A. Hillman: *Four-manifolds, Geometries and Knots.* Geometry & Topology Monographs, Vol. 5, 2002.

[20] F. Hirzebruch: The signature of ramified coverings, in *Global analysis: Papers in honor of K. Kodaira*, 253–266. University of Tokyo Press, Tokyo, 1969.

[21] S. T. Hu: *Homotopy Theory.* Academic Press Inc., Montreal, 1959.

[22] D. Husemoller: *Fibre Bundles.* Graduate Texts in Mathematics, No. 20, Springer-Verlag, New York-Heidelberg, 1966.

[23] C. Labruère and L. Paris: Presentations for the punctured mapping class groups in terms of Artin groups, *Algebr. Geom. Topol.* **1** (2001), 73–114.

[24] S. K. Lando and A. K. Zvonkin: *Graphs on Surfaces and Their Applications.* Encyclopaedia of Mathematical Sciences, Vol. 141, Springer-Verlag, New York, 2004.

[25] J. D. McCarthy: Private communication.

[26] D. McDuff and D. Salamon: *Introduction to Symplectic Topology.* Oxford Mathematical Monographs, Clarendon Press, Oxford, 1995.

[27] J. W. Morgan, Z. Szabó and C.H. Taubes: A product formula for the Seiberg-Witten invariants and the generalized Thom conjecture, *J. Differential Geom.* **44** (1996), 706–788.

[28] P. Ozsváth and Z. Szabó: The symplectic Thom conjecture, *Ann. Math.* **151** (2000), 93–124.

[29] L. Paris and D. Rolfsen: Geometric subgroups of mapping class groups, *J. Reine Angew. Math.* **521** (2000), 47–83.

[30] B. D. Park, M. Poddar and S. Vidussi: Homologous non-isotopic symplectic surfaces of higher genus, *Trans. Amer. Math. Soc.* (to appear).

[31] A. Scorpan: *The Wild World of* 4*-Manifolds.* Amer. Math. Soc., Providence, 2005.

[32] B. Siebert and G. Tian: On the holomorphicity of genus two Lefschetz fibrations, *Ann. of Math.* **161** (2005), 955–1016.

[33] I. Smith: Symplectic submanifolds from surface fibrations, *Pacific J. Math.* **198** (2001), 197–205.

[34] N. Steenrod: *The Topology of Fibre Bundles.* Princeton University Press, Princeton, 1974.

Index