

Durham E-Theses

Pseudo-isotopies and embedded surfaces in 4-manifolds

OLIVER SUNIL SINGH

How to cite:

SINGH, OLIVER SUNIL (2022) Pseudo-isotopies and embedded surfaces in 4-manifolds. Doctoral thesis, Durham University.

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a <https://etheses.durham.ac.uk/id/eprint/14601/> is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

Pseudo-isotopies and embedded surfaces in 4-manifolds

Oliver Singh

A thesis presented for the degree of
Doctor of Philosophy



Department of Mathematical Sciences
Durham University
United Kingdom

August 2022

Pseudo-isotopies and embedded surfaces in 4-manifolds

Oliver Singh

Submitted for the degree of Doctor of Philosophy

August 2022

Abstract: The focus of this thesis is the study of smooth 4-dimensional manifolds. We examine two problems relating to 4-manifolds, the first pertaining to pseudo-isotopies and diffeomorphisms of 4-manifolds, and the second pertaining to embedded surfaces in 4-manifolds. We summarise our key results below.

A diffeomorphism f of a compact manifold X is pseudo-isotopic to the identity if there is a diffeomorphism F of $X \times I$ which restricts to f on $X \times 1$, and which restricts to the identity on $X \times 0$ and $\partial X \times I$. We construct examples of diffeomorphisms of 4-manifolds which are pseudo-isotopic but not isotopic to the identity. To do so, we further understanding of which elements of the “second pseudo-isotopy obstruction”, defined by Hatcher and Wagoner, can be realised by pseudo-isotopies of 4-manifolds. We also prove that all elements of the first and second pseudo-isotopy obstructions can be realised after connected sums with copies of $S^2 \times S^2$.

If Σ and Σ' are homotopic embedded surfaces in a 4-manifold then they may be related by a regular homotopy (at the expense of introducing double points) or by a sequence of stabilisations and destabilisations (at the expense of adding genus). This naturally gives rise to two integer-valued notions of distance between the embeddings: the singularity distance $d_{\text{sing}}(\Sigma, \Sigma')$ and the stabilisation distance $d_{\text{st}}(\Sigma, \Sigma')$. We use techniques similar to those used by Gabai in his proof of the 4-dimensional light-bulb theorem, to prove that $d_{\text{st}}(\Sigma, \Sigma') \leq d_{\text{sing}}(\Sigma, \Sigma') + 1$.

Declaration

The work in this thesis is based on research carried out in the Department of Mathematical Sciences at Durham University. No part of this thesis has been submitted elsewhere for any degree or qualification.

Copyright © 2022 Oliver Singh.

“The copyright of this thesis rests with the author. No quotation from it should be published without the author’s prior written consent and information derived from it should be acknowledged.”

Acknowledgements

First and foremost, I would like to thank my supervisors Andrew Lobb and Mark Powell for their guidance, their patience, for the mathematics they taught me, and for their endless humour. Thanks to you both for encouraging me to spend three months in Bonn and for enduring my other ventures away from Durham.

I would like to thank all the people who made Durham maths department and its surrounding pubs an enjoyable place to be over the years, particularly Sam, Scott, John, and Anthony. Thanks for all the wisdom you imparted on topology, academia, set theory, and life.

Thanks to the Max Planck Institute for Mathematics for hosting me for three months while I worked on some of the contents of this thesis, and to all the people who made me feel welcome in the department and who kept me company during pandemic restrictions in Bonn.

Thanks to Ben for encouraging me to be a useful mathematician during Covid-19 and to the friends I made in the Cabinet Office for staying friends with me when I went back to being a less useful mathematician.

I am, of course, incredibly grateful to all the friends and family who kept me sane over the last 4 years. Thanks to Netti for keeping me sane in Durham and for your exceptional counsel, to Anne, Ben, Rachel and Robbie for keeping me company and providing hilarious distraction while I wrote up, to Shivani for letting me work in her flat so many times, and to you all for being great pals.

Finally, thanks to the Engineering and Physical Sciences Research Council (EPSRC) for supporting my PhD via Grant no. EP/N509462/1 project no. 1918079, awarded through Durham University.

Contents

Abstract	iii
1 Introduction	3
1.1 Pseudo-isotopies and diffeomorphisms of 4-manifolds	3
1.2 Distances between surfaces in 4-manifolds	8
1.2.1 The topologically flat case	10
2 Pseudo-isotopies and diffeomorphisms of 4-manifolds	13
2.1 Background	13
2.1.1 Morse theory and the Morse chain complex	13
2.1.2 Homotopies of surfaces in 4-manifolds and Whitney discs	15
2.1.3 The $\mathbb{Z}[\pi_1 X]$ intersection number of spheres	20
2.1.4 Postnikov towers and k -invariants	21
2.2 Functional approach to pseudo-isotopies	22
2.2.1 Generic paths of functions in \mathcal{F}	23
2.3 Geometric picture and conventions in dimension 4	26
2.3.1 Handle slides	27
2.3.2 Intersections of spheres	28
2.4 Review of Hatcher and Wagoner's $\text{Wh}_2(\pi_1 X)$ invariant Σ	28
2.4.1 Algebra of Wh_2	28
2.4.2 Definition of Σ	30

2.4.3	Reduction to eyes and $\ker \Sigma$	32
2.5	Stable surjectivity of Σ in dimension 4	32
2.5.1	Proof of stable surjectivity in dimension 4	35
2.6	The $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$ invariant Θ	38
2.6.1	Definition of $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$	39
2.6.2	Construction of Θ	40
2.6.3	Geometric description of 1-parameter families in $\ker \Sigma$ in dimension 4	44
2.6.4	Calculation for geometrically simple families	46
2.7	Realisation theorem for $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$	50
2.8	Pseudo-isotopy versus isotopy	61
2.8.1	Duality formulae	63
2.8.2	Inertial pseudo-isotopies of $M^3 \times I$	67
2.8.3	Diffeomorphisms of $X^4 \times I$	68
2.8.4	Diffeomorphisms of $S^1 \times S^2 \times I$	71
2.8.5	Diffeomorphisms of the connect sum of aspherical 3-manifolds times I	72
3	Distances between surfaces in 4-manifolds	83
3.1	Plan of the chapter.	84
3.2	Surface tubing diagrams	85
3.2.1	Tubed surfaces	86
3.3	The Tube Swap Lemma	88
3.4	The Tube Move Lemma	95
3.5	Proof of Theorem G	98
3.6	Destabilising the slice surfaces K and K'	107
3.6.1	Banded link presentations of knotted surfaces	107
3.6.2	Proof that K and K' destabilise to the standard disc	108

Chapter 1

Introduction

We study two topics in low dimensional topology, both relating to 4-dimensional manifolds. The first pertains to pseudo-isotopies of 4-manifolds, and the second pertains to embedded surfaces in 4-manifolds. The material here is based on our papers [Sin21] and [Sin20]. In this chapter we will outline our main results, in Chapter 2 we prove the results relating to pseudo-isotopies, and in Chapter 3 we prove the results relating to embedded surfaces.

1.1 Pseudo-isotopies and diffeomorphisms of 4-manifolds

Let X be a smooth compact manifold of dimension n . A *pseudo-isotopy* of X is a diffeomorphism $F: X \times [0, 1] \rightarrow X \times [0, 1]$ such that $F|_{X \times 0}$ and $F|_{\partial X \times [0, 1]}$ are the identity. We let $\mathcal{P} = \mathcal{P}(X, \partial X)$ be the *space of pseudo-isotopies of X* , which is a topological space when equipped with the C^∞ topology. When X is closed we write $\mathcal{P}(X) = \mathcal{P}(X, \emptyset)$.

Given diffeomorphisms $f, g: X \rightarrow X$ we say f and g are *pseudo-isotopic* if there exists a pseudo-isotopy of X such that $F|_{X \times 1}$ is $g^{-1} \circ f$. Thinking of an isotopy of X as a map $F: X \times [0, 1] \rightarrow X \times [0, 1]$ which is level preserving, that is $F(X \times t) = X \times t$ for every $t \in [0, 1]$, it is clear that if diffeomorphisms f and g of X are isotopic, then they are pseudo-isotopic.

The first aim of Chapter 2 is to extend the list of 4-manifolds for which the converse is known to be false. Denoting the subgroup of diffeomorphisms fixing the boundary which are pseudo-isotopic to the identity by $\text{Diff}_{PI}(X, \partial X)$, we construct non-trivial elements of $\pi_0 \text{Diff}_{PI}(X, \partial X)$ for certain 4-manifolds X .

Theorem A. *Let X be either the 4-manifold $S^1 \times S^2 \times I$ or $(M_1 \# M_2) \times I$, for M_1, M_2 closed, orientable, aspherical 3-manifolds. Then there is a subgroup $K \leq \pi_0 \text{Diff}_{PI}(X, \partial X)$ and a surjective map*

$$\Theta': K \longrightarrow \bigoplus_{i \in \mathbb{N}} \mathbb{Z}.$$

Hence there are infinitely many distinct isotopy classes of diffeomorphisms of X fixing the boundary, which are pseudo-isotopic to the identity.

Igusa points out in [Igu21b] the map $\pi_0 \text{Diff}_{PI}(M^{n-1} \times I, \partial(M^{n-1} \times I)) \rightarrow \pi_0 \text{Diff}_{PI}(M^{n-1} \times S^1)$ induced by gluing the top and bottom of $M \times I$ together is injective [Igu21b, Lemma 5.1].

This allows us to state the below corollary.

Corollary B. *Let X be either the 4-manifold $S^1 \times S^2 \times S^1$ or $(M_1 \# M_2) \times S^1$, for M_1, M_2 closed, orientable, aspherical 3-manifolds. Then there is a subgroup $K \leq \pi_0 \text{Diff}_{PI}(X)$ and a surjective map $\Theta': K \rightarrow \bigoplus_{i \in \mathbb{N}} \mathbb{Z}$. Hence there are infinitely many distinct isotopy classes of diffeomorphisms of X which are pseudo-isotopic to the identity.*

To prove Theorem A, we utilise the so called “second pseudo-isotopy obstruction” Θ , taking values in $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$, which was defined by Hatcher and Wagoner in [HW73] and refined by Igusa [Igu84]. In dimensions $n \geq 5$, Hatcher [Hat73] uses Θ to construct non-trivial elements of $\pi_0 \text{Diff}_{PI}(M^{n-1} \times I, \partial(M^{n-1} \times I))$. To extend this result to 4 dimensions we further understanding of which elements of $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$ are realised by Θ for 4-manifolds. This is the second aim of Chapter 2 .

Our third aim is to show that both Θ and the “first pseudo-isotopy obstruction”

$$\Sigma: \pi_0 \mathcal{P}(X, \partial X) \rightarrow \text{Wh}_2(\pi_1 X)$$

are (in some sense) surjective in dimension 4 after taking connect sums of X with copies of $S^2 \times S^2$. Both invariants are surjective without any such stabilisation in dimension ≥ 5 .

Before we state the rest of our main theorems, we recall some high-dimensional background. In high dimensions, pseudo-isotopies are classified up to isotopy by Θ and Σ . Building on work of Hatcher and Wagoner in [HW73] and [Hat73], Igusa shows in [Igu84] that for X a smooth manifold of dimension at least 6 there is a natural exact sequence

$$K_3 \mathbb{Z}[\pi_1 X] \xrightarrow{\chi} \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) \rightarrow \pi_0 \mathcal{P}(X, \partial X) \xrightarrow{\Sigma} \text{Wh}_2(\pi_1 X) \rightarrow 0.$$

Igusa shows that if the first Postnikov invariant k_1X vanishes, then χ is 0 and the sequence splits, with splitting

$$\Theta_\sigma: \pi_0\mathcal{P} \rightarrow \text{Wh}_1(\pi_1X, \mathbb{Z}_2 \times \pi_2X)$$

dependent on a choice of section $\sigma: X_{(1)} \rightarrow X_{(2)}$, where $X_{(i)}$ is the i th stage in a Postnikov tower for X . The restriction of this map to $\ker \Sigma$

$$\Theta: \pi_0 \ker \Sigma \rightarrow \text{Wh}_1(\pi_1X, \mathbb{Z}_2 \times \pi_2X),$$

originally defined by Hatcher and Wagoner in [HW73], does not depend on a choice of section. When $k_1X \neq 0$ it follows from the constructions in [Igu84] that there is a map

$$\Theta: \ker \Sigma \rightarrow \text{Wh}_1(\pi_1X; \mathbb{Z}_2 \times \pi_2X) / \chi(K_3\mathbb{Z}[\pi_1X]).$$

The maps Σ and Θ (restricted to $\ker \Sigma$) are also defined in 4 and 5 dimensions. In dimension $n \geq 5$ Hatcher and Wagoner show that Σ and Θ are surjective. However, they are not able to prove this in dimension 4, where the situation is less clear.

Convention. Throughout, 4-manifolds will be smooth, compact, and connected.

We prove that in dimension 4, the following elements of the second obstruction group are realised.

Theorem C. *For X a compact 4-manifold, let*

$$\begin{aligned} \Xi &= \left\langle (s + \sigma)\gamma \mid w_2^X(\sigma) \neq 0 \text{ or } s = 0, s \in \mathbb{Z}_2, \sigma \in \pi_2X, \gamma \in \pi_1X \right\rangle \\ &\subset (\mathbb{Z}_2 \times \pi_2X)[\pi_1X] / \langle \alpha\gamma - \alpha^\tau \tau\gamma\tau^{-1}, \beta \cdot 1 \mid \alpha, \beta \in \mathbb{Z}_2 \times \pi_2X, \tau, \gamma \in \pi_1X \rangle \\ &= \text{Wh}_1(\pi_1X; \mathbb{Z}_2 \times \pi_2X). \end{aligned}$$

If $k_1X = 0$ then $\Xi \subset \Theta(\ker \Sigma)$. Otherwise the same is true passing to the quotient $\text{Wh}_1(\pi_1X; \mathbb{Z}_2 \times \pi_2X) / \chi(K_3\mathbb{Z}[\pi_1X])$.

For the identification of $\text{Wh}_1(\pi_1X; \mathbb{Z}_2 \times \pi_2X)$ with

$$(\mathbb{Z}_2 \times \pi_2X)[\pi_1X] / \langle \alpha\sigma - \alpha^\tau \tau\sigma\tau^{-1}, \beta \cdot 1 \mid \alpha, \beta \in \mathbb{Z}_2 \times \pi_2X, \tau, \sigma \in \pi_1X \rangle$$

see Corollary 2.6.4. The proof of Theorem C involves a detailed analysis of the Θ obstruction in terms of Whitney discs and framings of these discs in the 4-dimensional

“middle-middle level” of 1-parameter families of handlebody structures. We believe this to be of independent interest; see Propositions 2.6.8 and 2.6.10.

Jahren [Jah], in an unpublished work, proves a similar theorem by different methods. Specifically he proves that all elements of $\text{Wh}_1(\pi_1 X, \pi_2 X)$ are realised, and that when X contains an odd sphere, all elements of $\text{Wh}_1(\pi_1 X, \mathbb{Z}_2)$ are realised. We obtain both of these results as a corollary of Theorem C.

Corollary 1.1.1.

1. For any 4-manifold X , Θ surjects onto $\text{Wh}_1(\pi_1 X, \pi_2 X)$.
2. If X is a 4-manifold with an odd sphere, that is $S \in \pi_2 X$ with $S \cdot S$ odd (note that we do not require S to be embedded), then Θ surjects onto $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2)$.

Corollary 1.1.1 extends a recent result by Igusa in concurrent work [Igu21a, Theorem A, Theorem B]. Using different methods, Igusa proves Corollary 1.1.1.(2) with the additional requirement that S is embedded. Igusa also proves that certain elements of $\text{Wh}_1(\pi_1 X, \pi_2 X)$ are realised, namely ones where the element of $\pi_2 X$ can be represented by an embedded sphere.

We also prove that after stabilisation of X with a single $S^2 \times S^2$ we may realise any element of $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$.

Theorem D. *Let X be a compact 4-manifold. Note that $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$ includes in*

$$\text{Wh}_1(\pi_1(X \# S^2 \times S^2); \mathbb{Z}_2 \times \pi_2(X \# S^2 \times S^2)),$$

and identify $x \in \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$ with its image under this inclusion. There is a pseudo-isotopy F of $X \# S^2 \times S^2$, which is in $\ker \Sigma$ such that

$$\Theta(F) = x \in \text{Wh}_1(\pi_1(X \# S^2 \times S^2); \mathbb{Z}_2 \times \pi_2(X \# S^2 \times S^2)) / \chi(K_3 \mathbb{Z}[\pi_1(X \# S^2 \times S^2)]).$$

We also prove a stable version of Hatcher and Wagoner’s surjectivity result for Σ .

Theorem E. *Let X be a compact 4-manifold and $x \in \text{Wh}_2(\pi_1 X)$. There exists N , and a pseudo-isotopy F of $X \#^N S^2 \times S^2$ such that*

$$\Sigma(F) = x \in \text{Wh}_2(\pi_1(X \#^N S^2 \times S^2)) = \text{Wh}_2(\pi_1 X).$$

As a consequence of Theorem C we are also able to construct diffeomorphisms of certain 5-manifolds.

Theorem F. *Suppose X is a 4-manifold which contains an element $\sigma \in \pi_2(X)$ with $w_2^X(\sigma) \neq 0$, and an element $\gamma \in \pi_1 X$ such that γ and γ^{-1} are not conjugate, and suppose also that either $k_1 X = 0$ or $K_3 \mathbb{Z}[\pi_1 X] = 0$. Then in $\text{Diff}(X \times I, \partial(X \times I))$ there exist diffeomorphisms pseudo-isotopic to the identity but not isotopic to it.*

This result is analogous to a result by Hatcher [Hat73, Corollary 4.5], which gives diffeomorphisms of manifolds of dimension greater than 6.

The other examples of diffeomorphisms of 4-manifolds which are pseudo-isotopic but not isotopic to the identity that we know of come from Budney-Gabai [BG21], Watanabe [Wat20], Igusa [Igu21b], and the following examples from gauge theory: Ruberman [Rub98], [Rub99], Baraglia-Konno [BK20], Kronheimer-Mrowka [KM20], and Lin [Lin20]. Budney and Gabai construct diffeomorphisms of $S^1 \times B^3$ and $S^1 \times S^3$, and by related methods Watanabe constructs diffeomorphisms of $\Sigma(2, 3, 5) \times S^1$. Budney and Gabai's examples are pseudo-isotopic to the identity by Sato [Sat69] and Lashof-Shaneson [LS69], while Watanabe's examples are pseudo-isotopic to the identity by [Wat20, Theorem 1.8 and Theorem 9.3], which Watanabe attributes to Teichner. Igusa in a work concurrent to ours constructs diffeomorphisms of $((S^1 \times S^2) \# M) \times S^1$ for M a non-simply connected 3-manifold. The gauge theoretic examples on the other hand give diffeomorphisms of simply connected 4-manifolds which are not isotopic to the identity but induce the identity on homology; by Kreck [Kre79, Theorem 1], these diffeomorphisms are pseudo-isotopic to the identity.

Other work that has been done on pseudo-isotopies of 4-manifolds includes the work of Quinn [Qui86] who proves that topological pseudo-isotopies of simply connected 4-manifolds are topologically isotopic rel boundary to an isotopy. Hence in the topological setting isotopy and pseudo-isotopy are the same for simply connected 4-manifolds.

Kwasik [Kwa87] shows that the same is not true for non simply connected 4-manifolds, and asserts that if topological pseudo-isotopies are allowed then Θ is surjective. However, it is unclear to us how to define Θ and Σ in the topological setting, and unclear how to show it is well defined, particularly as current definitions are heavily reliant on Cerf theory.

It appears to us that a topological definition of both Θ and Σ would be of value, and could facilitate further results in the topological setting.

1.2 Distances between surfaces in 4-manifolds

Let X be a smooth, compact, orientable 4-manifold, possibly with boundary. Let Σ, Σ' be connected, oriented, compact, smooth, properly embedded surfaces in X . We say that Σ' is a *stabilisation* of Σ if there is an embedded solid tube $D^1 \times D^2 \subset X$ such that $\Sigma \cap (D^1 \times D^2) = \{0, 1\} \times D^2$, and Σ' is obtained from Σ by removing these two discs and replacing them with $D^1 \times S^1$, as in Figure 1.1, and then smoothing corners. In this situation we say that Σ is a *destabilisation* of Σ' .




Figure 1.1: A stabilisation. Given $D^1 \times D^2 \subset X$ which intersects Σ on $S^0 \times D^2$, we remove the two discs $S^0 \times D^2$, add the tube $D^1 \times S^1$, then smooth corners.

Definition 1.2.1. Given Σ, Σ' as above, both of genus g , define the *stabilisation distance* between Σ and Σ' to be

$$d_{\text{st}}(\Sigma, \Sigma') = \min_{\mathbb{S}} \max\{|g(P_1) - g|, |g(P_2) - g| \dots, |g(P_k) - g|\},$$

where \mathbb{S} is the set of sequences P_1, \dots, P_k of connected, oriented, embedded surfaces where $\Sigma = P_1, \Sigma' = P_k$ and P_{i+1} differs from P_i by one of, i) stabilisation, ii) destabilisation, or iii) ambient isotopy. If no such sequence exists we declare $d_{\text{st}}(\Sigma, \Sigma') = \infty$.

Definition 1.2.2. Given an immersion $f: S \rightarrow X$, define $\text{sing}(f) \subset X$ to be the set of double points. Define the *singularity distance* to be

$$d_{\text{sing}}(\Sigma, \Sigma') = \frac{1}{2} \min_H \max_{t \in [0,1]} |\text{sing}(H_t)|,$$

where the minimum is taken over all generic regular homotopies H from Σ to Σ' . If no such regular homotopy exists we declare $d_{\text{sing}}(\Sigma, \Sigma') = \infty$.

Remark 1.2.3. It is possible that $d_{\text{st}}(\Sigma, \Sigma') < \infty$ and $d_{\text{sing}}(\Sigma, \Sigma') = \infty$. For example let $U \subset S^3$ be the unknot, let $\Sigma = U \times S^1 \subset S^3 \times S^1$, and let Σ' be a trivially embedded torus in $S^3 \times \{\text{pt}\} \subset S^3 \times S^1$. Since Σ and Σ' are not homotopic $d_{\text{sing}}(\Sigma, \Sigma') = \infty$. Further Σ and Σ' are related by a destabilisation then a stabilisation, so $d_{\text{st}}(\Sigma, \Sigma') = 1$.

The main theorem of Chapter 3 is the following:

Theorem G. *If $\Sigma, \Sigma' \subset X$ are connected, smooth, properly embedded, oriented surfaces of the same genus then*

$$d_{\text{st}}(\Sigma, \Sigma') \leq d_{\text{sing}}(\Sigma, \Sigma') + 1.$$

The proof is constructive, in that given a regular homotopy with at most $2n$ double points at each time, we construct a sequence of stabilisations and destabilisations such that the maximum genus occurring is at most $g + n + 1$. Since embedded oriented surfaces are regularly homotopic if and only if they are homotopic, see Remark 3.0.1, this gives a new proof of a certain case of Baykur and Sunukjian's result [BS15, Theorem 1].

Corollary H. *If $\Sigma, \Sigma' \subset X$ are homotopic, connected, smooth, properly embedded, oriented surfaces, then they differ by a sequence of stabilisations, destabilisations, and ambient isotopy.*

We do not know whether the $+1$ in Theorem G is really essential. As we discuss in Remark 3.5.2(2), it does not seem obvious how to adapt the proof to prove the inequality without the $+1$. It would, however, be interesting to find an example where the sequence of stabilisations constructed in the proof is minimal.

Conjecture 1.2.4. *There exists a smooth, compact, orientable 4-manifold X and smooth, properly embedded, orientable, regularly homotopic surfaces Σ and Σ' with*

$$d_{\text{st}}(\Sigma, \Sigma') = d_{\text{sing}}(\Sigma, \Sigma') + 1.$$

Juhász and Zemke define invariants in [JZ18] which bound below the minimum of two similarly defined distances, μ_{st} and μ_{sing} using their notation. Their singularity distance μ_{sing} is the same as d_{sing} defined here, however $\mu_{\text{st}} \leq d_{\text{st}}$, as their definition of stabilisations and destabilisations additionally allows taking connect sum with a 2-knot (or removing a

2-knot attached via connect sum). We currently do not know of any invariants that can bound d_{st} below without also bounding d_{sing} below.

The techniques we use to turn regular homotopies into sequences of stabilisations are inspired by those used by Gabai [Gab17] to prove that if two homotopic embedded orientable surfaces have a common transverse sphere in X , where $\pi_1(X)$ has no 2-torsion, then the surfaces are ambiently isotopic (and so both the distances defined here are 0). A transverse sphere for a given embedded sphere $S \subset X$ is another embedded sphere $P \subset X$ such that P intersects S transversely at a single point.

Miller [Mil19] also recently proved, under the assumption that one of the surfaces has a transverse sphere, and that $\pi_1(X)$ has no 2-torsion, that regularly homotopic surfaces are concordant. There are also modified statements of these theorems when $\pi_1(X)$ has 2-torsion.

Schwartz found infinitely many examples [Sch18] which demonstrate that the assumption on $\pi_1(X)$ was essential. She produced 4-manifolds X with $\pi_1(X) \cong \mathbb{Z}_2$, and pairs of homotopic embedded 2-spheres in X which share a transverse sphere but are not concordant. We would be interested to know the stabilisation and double point distances of these examples.

Schneiderman and Teichner [ST19] recently reproved the 4-dimensional light bulb theorem and characterised the situation for general $\pi_1(X)$ using an obstruction of Freedman and Quinn.

1.2.1 The topologically flat case

The condition of smoothness above was required to define regular homotopy. We define regular homotopy in a topological 4-manifold X by saying two locally flat surfaces $\Sigma, \Sigma' \subset X$ are *regularly homotopic* if they differ by a sequence of finger moves, Whitney moves (where the Whitney discs are locally flatly embedded), and ambient isotopy.

We also define stabilisation in the same way, dropping the smoothness condition on the embedding of $D^1 \times D^2$. Using the topological definitions, we define $d_{\text{st}}^{\text{top}}$ and $d_{\text{sing}}^{\text{top}}$ analogously. Repeating the proof of Theorem G, without the initial smoothness assumption yields:

Theorem I. *If Σ and Σ' are orientable, compact, connected, locally flat, properly embedded surfaces in X of the same genus, then,*

$$d_{\text{st}}^{\text{top}}(\Sigma, \Sigma') \leq d_{\text{sing}}^{\text{top}}(\Sigma, \Sigma') + 1.$$

As a consequence, we prove an analogue of Corollary H for locally flat surfaces.

Corollary J. *Let Σ and Σ' be orientable, compact, connected, locally flat, properly embedded surfaces in X . If Σ and Σ' are regularly homotopic (topologically), then they differ by a sequence of topological stabilisations, destabilisations, and ambient isotopy.*

Chapter 2

Pseudo-isotopies and diffeomorphisms of 4-manifolds

In this chapter we address pseudo-isotopies of 4-manifolds and prove the results outlined in Section 1.1.

2.1 Background

We first recall some general topology background. We will review some standard tools of Morse theory, we will also review Whitney discs and operations on Whitney discs, as well as Postnikov towers and k -invariants.

As is standard in Morse theory and Cerf theory, we fix a Riemannian metric μ on X and take the product metric on $X \times I$.

Remark 2.1.1. The choice of metric is not important, and in fact it is possible to avoid fixing a metric entirely. Indeed Hatcher and Wagoner consider a space $\widehat{\mathcal{F}}(X, \partial X)$ which includes all possible choices of metric, and show that it is homotopy equivalent to the space $\mathcal{F}(X, \partial X)$ which we define in Section 2.2; see [HW73, Chapter 1, §3].

2.1.1 Morse theory and the Morse chain complex

Let f be a Morse function

$$f: X \times I \rightarrow I$$

such that $f|_{\partial X \times I}$ is the projection to I and that $f(X \times i) = i$ for $i \in \{0, 1\}$. Suppose also that f has no critical points near $X \times \partial I$ or $\partial X \times I$. By choice of Riemannian metric we have a uniquely defined gradient ∇f . Given $s \in \mathbb{R}$ we denote the *flow* of f by $\phi_{f,s}: X \times I \rightarrow X \times I$, where $\phi_{f,s}(p)$ is given by pushing p along ∇f for time s .

Recall that the *trajectory* of a point p is given by $\{\phi_{f,s}(p) | s \in \mathbb{R}\}$, and that trajectories are embedded copies of $(0, 1)$, $(0, 1]$, $[0, 1)$ or $[0, 1]$, with any boundary points lying in $X \times \partial I$. Recall also that the limit point on approaching the open end of a trajectory is always a critical point. Given a critical point p we denote the *stable set* by

$$W_f(p) = \{q \mid \lim_{s \rightarrow +\infty} \phi_{f,s}(q) = p\}$$

and the *unstable set* by

$$W_f^*(p) = \{q \mid \lim_{s \rightarrow -\infty} \phi_{f,s}(q) = p\}.$$

Given two critical points $p, q \in X \times I$ of index i and j respectively, let

$$T_p^q = \left\{ \{\phi_{f,s}(a)\}_{s \in \mathbb{R}} \mid a \in W(p) \cap W^*(q) \right\}$$

be the set of trajectories from q to p . We refer to these trajectories as *j/i trajectories* (the j being on top of the fraction indicates that is also “on top” in the manifold, since a j/i trajectory goes down from a critical point of index j to one of index i). If p, q are critical points of index i and $i + 1$ respectively and f is a Morse function in general position, a dimension count shows that $W(p) \cap W^*(q)$ is a collection of isolated arcs and so there are finitely many trajectories from q to p .

The Morse chain complex

We can capture the data of the $\frac{i}{i+1}$ intersections in a chain complex, the *Morse chain complex*, defined as follows. Let $\psi: \tilde{X} \rightarrow X$ be the universal cover of X , and let f be a self indexing Morse function on X , i.e. there exists $r_0 = 0 < r_1 < \dots < r_{n+1} = 1 \in [0, 1]$ with all critical points of index i having critical values between r_i and r_{i+1} . Let $V_i = f^{-1}([r_i, r_{i+1}])$, and let $\tilde{V}_i = \psi^{-1}(V_i)$. Define the chain complex by

$$C_i(X) = H_i(\tilde{V}_i, \tilde{V}_{i-1})$$

which is a finitely generated free module over $\mathbb{Z}[\pi_1(X)]$ with a generator for each critical point of index i .

Fix a base point $p \in X$ and a lift of the base point $\tilde{p} \in \tilde{X}$, and for each critical point c of f pick a path γ_c from p to c . This gives a basis for C_i as a finitely generated free $\mathbb{Z}[\pi_1(X)]$ module. We can describe the differential as follows. Let S_i be the collection of index i critical points, recall that T_b^a is the set of trajectories from a to b . Using the basis above, this determines a matrix $(\partial_i^f : C_i \rightarrow C_{i-1})$ given by

$$\partial_i^f(c) = \sum_{b \in S_{i-1}} \sum_{\psi \in T_b^c} [\gamma_c \cdot \psi \cdot \gamma_b^{-1}]b.$$

2.1.2 Homotopies of surfaces in 4-manifolds and Whitney discs

We will wish to study deformations of immersed surfaces in 4-manifolds, and so recall some standard techniques for doing so; for full details see [FQ90].

Let X be a smooth 4-manifold. Given a compact immersed surface $S \subset X$ (with possibly many components), recall that a *regular homotopy* of this surface in X is a homotopy through immersions; i.e. a smooth map $H : \Sigma \times I \rightarrow X$ such that $H(-, 0)$ maps Σ to the surface S , and $H(-, t)$ is an immersion for all t . We frequently omit the map and discuss only the immersed surface, which we denote $S_t = H(\Sigma, t)$.

Given an immersed surface, by standard general position arguments we can perturb it so that it has only finitely many isolated self-intersections, and that these self-intersections are transverse double points; that is the immersion is 2 to 1 at these points (as opposed to n to 1 for $n > 2$).

By general position we may perturb a regular homotopy so that for each t , S_t has only transverse double-points, except at finitely many values of t . At these values of t we see a non-transverse point; near these values, the homotopy can always be described as one of the following local moves; see [FQ90] for further details.

Definition 2.1.2. Let $S \subset X$ be an immersed oriented surface, and let γ be an arc such that the endpoints are in S , but the interior is disjoint from S . Then we may perform a *finger move*, pushing one sheet of the surface along this arc into the other to introduce two points of transverse intersection; see the first image in Figure 2.1.

Convention. Throughout this chapter it will be useful to us to insist that the finger move arc γ in Definition 2.1.2 be an oriented arc (one may freely pick the orientation, we just require that it is kept track of).




Figure 2.1: The time slice where the finger moves and Whitney moves occur. Note that the horizontal line continues into the past and future, and is unchanged by the homotopies.

Remark 2.1.3. Let S be a collection of spheres embedded in X , and suppose for each sphere we have an arc to a basepoint $* \in X$. Then we can associate an element of $\pi_1(X)$ to any finger move arc by going along a basepoint arc, along some arc in S to the base of the finger move arc, along the finger move arc, along another arc in S to the basepoint arc, then backwards along another basepoint arc (possibly the same basepoint arc if the endpoints of the finger move arc lie on the same component of S). We call this the element of $\pi_1(X)$ associated to the finger move. If additionally S is π_1 -negligible (that is $\pi_1(X) = \pi_1(X \setminus S)$) then the finger move is uniquely determined by its associated element of $\pi_1(X)$; indeed for any $\gamma \in \pi_1(X)$ there always exists some choice of finger move arc with associated element γ and any two choices can easily be shown to be homotopic in X , and hence in $X \setminus S$ since S is π_1 -negligible.

Remark 2.1.4. Note that reversing the orientation of γ results in an isotopic finger move, and inverts the associated element of $\pi_1 X$. Later, the surfaces between which we perform finger moves will have a natural order, so it will be useful to consider a finger move “from” one surface “to” another.

Remark 2.1.5. Given an oriented immersed surface $S \subset X$ we wish to assign a sign to the double points of S . When X is oriented this is easy, however when X is non-orientable

we must work a little harder. In this case, we restrict to the case that S is π_1 -trivial; that is $\pi_1 S$ includes as 0 in $\pi_1 X$. Let $S = \cup_i S_i$, where S_i are the connected components. We pick a basepoint $* \in X$ and arcs from $*$ to each component S_i . We pick a local orientation at $*$ then transport the orientation at $*$ along the arcs to give an orientation of the total space of $\nu(S_i, X)$ for each i . This allows us to define signs of intersection as for an oriented manifold. Note that the signs of the intersections depend on our choices of arcs, but that for any two points $p, q \in S_i \cap S_j$ the dichotomy of these having the same sign or a different sign does not depend on our choices.

We now recall the definition of a Whitney disc, and a Whitney move.

Definition 2.1.6. Let $S \subset X$ be an oriented immersed surface (and that S is π_1 -trivial if X is non-orientable). Suppose z^+, z^- are double points of S with sign $+$ and $-$ respectively. Suppose $\alpha, \beta \subset S$ are embedded arcs, both of which have endpoints z^+ and z^- , and which are disjoint except at these endpoints. Suppose also that α and β do not intersect any intersection points of S except at their endpoints. We also require that the endpoints of α live in different sheets of S to the endpoints of β of S ; that is, letting S be the image of some immersion $f: \Sigma \rightarrow X$ we require that $f^{-1}(\alpha)$ and $f^{-1}(\beta)$ are disjoint including at the endpoints.

If $W \subset X$ is an immersed disc with boundary $\alpha \cup \beta$ we call W a *Whitney disc*. We require that W meets S transversely on ∂W . We call α and β the *Whitney arcs*.

Definition 2.1.7. Given a Whitney disc $W \subset X$, we pick an orientation for W . The bundle $\nu(W, X)$ is a 2-dimensional bundle over a disc and so is trivial and has a unique trivialisaton determined by the orientation. This trivialisaton induces a trivialisaton, or *framing*, of $\nu(W, X)|_{\partial W} \cong S^1 \times D^2$ which we call the *disc framing*.

There is another framing of $\nu(W, X)|_{\partial W}$ defined as follows. We define a section of $\nu(W, X)|_{\partial W}$, denoted $s: \partial W \rightarrow \nu(W, X)|_{\partial W}$, by requiring that s be parallel to S along β , and normal to S along α ; note that W is transverse to S on ∂W so this uniquely determines s up to homotopy. We call s the *Whitney section*. Now s is a section of a 2-dimensional bundle over S^1 , and so uniquely determines a trivialisaton (up to orientation, we use the orientation on $\nu(W, X)|_{\partial W}$ induced by the orientation of W and the orientation of X when X is oriented; when X is non orientable we use a local orientation as in Remark 2.1.5). We call this trivialisaton the *Whitney framing*.

If the Whitney framing and the disc framing agree (up to isotopy), we say that W is a *correctly framed Whitney disc*, or simply a *framed Whitney disc*.

Remark 2.1.8. If $\nu(W, X)|_{\partial W}$ is oriented, then if we fix a framing $\nu(W, X)|_{\partial W} \cong S^1 \times D^2$, then any other framing with the same orientation gives an element of $\pi_1 \mathrm{GL}_2(\mathbb{R}) = \mathbb{Z}$, and this element of \mathbb{Z} is canonical. Hence given two framings we have a well defined difference between them in \mathbb{Z} , using the first framing to identify $\nu(W, X)|_{\partial W} \cong S^1 \times D^2$ and the second to give the element of $\pi_1 \mathrm{GL}_2(\mathbb{Z}) = \mathbb{Z}$.

Remark 2.1.9. Equivalently and more succinctly one may say a Whitney disc is framed if the Whitney section extends to a section of $\nu(W, X)$. However, as we will perform moves to change the disc framing and the Whitney framing separately, it will be useful for us to consider both framings independently.

We now describe the second local move.

Definition 2.1.10. Given a correctly framed, embedded Whitney disc, as in Definition 2.1.7, whose interior is disjoint from S , we may perform a *Whitney move*, a regular homotopy removing the two double points z^+ and z^- ; see Figure 2.1.

Interior twists and boundary twists of discs

Suppose that we have an immersed Whitney disc W which pairs two intersections of some surface S (S is possibly disconnected, and W may intersect S away from the Whitney arcs). As in Remark 2.1.8 we denote the difference between the Whitney framing and the disc framing by $n_W \in \mathbb{Z}$. There are two operations to create a new Whitney disc described in [FQ90, Section 1.3]: the interior twist and the boundary twist. We recall briefly these operations and their effects on n_W below.

A *positive interior twist* is an operation which alters W in a small neighbourhood in the interior of W . In this neighbourhood we add an additional positive self-intersection to W . We call the resulting disc W' . Note that $\nu(W, X)|_{\partial W} = \nu(W', X)|_{\partial W'}$, and that the Whitney framing does not change. By perturbing W we can make $W' \cup -W$ immersed, one can easily see that $e(\nu(W' \cup -W, X)) = 2$, and so the disc framings of W and W' must differ by 2 (considering the clutching construction of bundles over spheres). Hence $n_{W'} = n_W + 2$.

Similarly a *negative interior twist* introduces an additional negative self-intersection and $n_{W'} = n_W - 2$.

A *positive boundary twist* alters W only in a neighbourhood of one Whitney arc; note that we may choose which arc. It introduces an additional single positive intersection between W and S by twisting W positively around S along this arc; see [FQ90, Section 1.3]. We call the result W' . In this case $n_{W'} = n_W + 1$.

Similarly a *negative boundary twist* introduces a additional single negative intersection between W' and S , and $n_{W'} = n_W - 1$.

We will frequently use these moves to obtain a correctly framed Whitney disc from one which is not correctly framed.

Pushing down

Let W be a Whitney disc in a 4-manifold V which pairs two intersections of some surface $A \subset X$. If W intersects some surface U in some point p , we may remove the intersection between W and U by *pushing down* the intersection into A . To do this we perform a finger move between U and A , resulting in two intersections between U and A ; see Figure 2.2.




Figure 2.2: A depiction of the pushing down operation, which turns an intersection between U and W into two intersections between U and A .

We can also push down to trade self intersections of W for intersections between W and A .

Transverse spheres and the Norman trick

Let A be a surface in a 4-manifold V , and let $A^* \subset V$ be a (possibly non-embedded) sphere with trivial normal bundle which intersects A transversely in a single point. We say A^* is a *transverse sphere* for A .

Given such a surface with a transverse sphere, if A intersects some other surface U in some point p , we may find a surface U' with one fewer intersection with A , by taking an embedded arc $\gamma \subset A$ from p to $A \cap A^*$, and tubing U to a parallel copy of A^* along the arc γ ; see Figure 2.3. This operation is called the *Norman trick*.




Figure 2.3: Using the transverse sphere A^* to perform the Norman trick, removing the intersection between A and U .

2.1.3 The $\mathbb{Z}[\pi_1 X]$ intersection number of spheres

Let A and B be immersed spheres in a 4-manifold which intersect transversely (or more generally any two connected π_1 -trivial surfaces). Suppose also that we have chosen a basepoint $* \in X$, and that we have chosen arcs $\alpha, \beta \subset X$ from $*$ to A and B respectively; denote the endpoints at which these arcs meet A and B by $*_A \in A$ and $*_B \in B$ respectively. Then given $p \in A \cap B$ we can assign a value in $\pm\pi_1 X$ to this intersection point as follows. The sign comes from the sign of the intersection; in the case that X is non-orientable this is as in Remark 2.1.5. The element of $\pi_1 X$ comes from taking an arc $a \subset A$ from $*_A$ to p and an arc $b \subset B$ from $*_B$ to p , then $\alpha \cdot a \cdot b^{-1} \cdot \beta^{-1}$ gives the element of $\pi_1 X$; see Figure 2.4.

We permit the arcs a and b may run over the intersection points and self-intersection points of A and B , but they may not change sheets of the surface at these double points; equivalently, one must be able to make a and b disjoint from the intersection and self-intersection points. Note that since A and B are π_1 -trivial $\alpha \cdot a \cdot b^{-1} \cdot \beta^{-1}$ is independent of the choice of a and b .

We can define an intersection number $A \cdot B \in \mathbb{Z}[\pi_1 X]$ by summing over all $p \in A \cap B$. Note that this intersection number depends on the choice of arcs α and β from A and B




Figure 2.4: The loop $\alpha \cdot a \cdot b^{-1} \cdot \beta^{-1}$.

to $*$; when we wish to consider intersection numbers in $\mathbb{Z}[\pi_1 X]$ we will fix some choice of arcs to the basepoint.

2.1.4 Postnikov towers and k -invariants

Since many of our theorems refer to the first Postnikov invariant (or k -invariant) $k_1 X$, we recall the basics of Postnikov towers and k -invariants below. For a more detailed treatment we direct the reader to [Whi78, pg. 421–437].


Recall that an *Eilenberg–MacLane space* $K(\pi, n)$ is a space whose n th homotopy group is isomorphic to π , with all other homotopy groups trivial.

Definition 2.1.11. Given a path connected space X , suppose we have spaces

$$X_{(0)}, X_{(1)}, \dots, X_{(n)}, \dots$$

and maps $p_n: X_{(n)} \rightarrow X_{(n-1)}$, $f_n: X \rightarrow X_{(n)}$, such that

1. the diagram below commutes.



2. The induced map $(f_n)_*: \pi_i(X) \rightarrow \pi_i(X_{(n)})$ is an isomorphism for $i \leq n$, and that $\pi_i(X_{(n)}) = 0$ for $i > n$.
3. The map $p_n: X_{(n)} \rightarrow X_{(n-1)}$ is a fibration, with fiber a $K(\pi_n(X), n)$ space.

We call such a tower of spaces and maps a *Postnikov system* (or *Postnikov tower*) and we call $X_{(n)}$ the Postnikov n -type of X .

Postnikov towers always exist for X a connected CW-complex, and uniquely determine X up to weak homotopy; see [Whi78, Chapter XI].

For each n consider the fibration

$$K(\pi_{n+1}X, n+1) \longrightarrow X_{(n+1)} \xrightarrow{p_{n+1}} X_{(n)}.$$

there is a well defined homology class $k_n \in H^{n+2}(X_{(n)}; \pi_{n+1}X)$ which classifies this fibration, see [Whi78, pg. 421–437]. In particular there exists a section of this fibration $X_{(n)} \rightarrow X_{(n+1)}$ if and only if $k_n = 0$. We call k_n the n th *k-invariant*. Note that our indexing is different to that of [Whi78], in order to agree with the indexing of Igusa [Igu84].

We are particularly interested in k_1X . Since $\pi_n(X_{(1)}) = 0$ for $n > 1$, in fact $X_{(1)}$ is itself an Eilenberg–MacLane space, $X_{(1)} = K(\pi_1X, 1)$, hence

$$k_1X \in H^3(K(\pi_1X, 1), \pi_2X) = H^3(\pi_1X; \pi_2X)$$

where $H^3(\pi_1X; \pi_2X)$ is group cohomology with coefficients twisted by the action of π_1X on π_2X .

2.2 Functional approach to pseudo-isotopies

We begin this section with a review of Cerf’s view of pseudo-isotopies as paths of Morse functions as set out by Hatcher and Wagoner in [HW73, Chapter 1, §2].

Let $I = [0, 1]$ and let $\mathcal{F} = \mathcal{F}(X, \partial X)$ be the space of C^∞ functions $f: X \times I \rightarrow I$ such that $f(x, 0) = 0$, $f(x, 1) = 1 \forall x$, and such that f has no critical points near $X \times 0$, $X \times 1$ or $\partial X \times I$. Let $\mathcal{E} \subset \mathcal{F}$ be the subset of all such functions with no critical points.

Denote the standard projection to I by $p: X \times I \rightarrow I$. We define a map

$$\begin{aligned}\Pi: \mathcal{P} &\longrightarrow \mathcal{F} \\ F &\longmapsto p \circ F\end{aligned}$$

Since any $F \in \mathcal{P}$ is a diffeomorphism, $p \circ F$ has no critical points so

$$\Pi(\mathcal{P}) \subset \mathcal{E}.$$

In fact $\Pi(\mathcal{P}) = \mathcal{E}$. Given $f \in \mathcal{E}$ we construct $F \in \mathcal{P}$ with $p \circ F = f$ by defining

$$F(x, s) = \phi_{f,s}(x, 0).$$

Further, Π is a fibration $\mathcal{I} \rightarrow \mathcal{P} \xrightarrow{\Pi} \mathcal{E}$ with fiber

$$\mathcal{I} = \{F: X \times I \rightarrow I \mid F|_{X \times 0} = \mathbf{1}_{X \times 0}, F(X \times t) = X \times t \forall t \in I\}$$

that is, the space of isotopies of X fixing $X \times 0$. The fiber \mathcal{I} is contractible via

$$\begin{aligned}H_s: \mathcal{I} &\rightarrow \mathcal{I} \\ H_s(F)(x, t) &= F(x, (1-s)t).\end{aligned}$$

Hence Π is a homotopy equivalence. Additionally \mathcal{F} is contractible, so we have

$$\pi_0 \mathcal{P} = \pi_0 \mathcal{E} = \pi_1(\mathcal{F}, \mathcal{E}).$$

In order to measure whether a given pseudo-isotopy F is isotopic to the identity, our strategy will be to join $\Pi(F)$ to p by a path $f_t \in \mathcal{F}$, and try to deform this path, fixing the ends, to lie in \mathcal{E} ; if we succeed then the path f_t is the trivial element of $\pi_1(\mathcal{P}, \mathcal{E})$, so F is the trivial element of $\pi_0 \mathcal{P}$. Conversely, our obstructions will be obstructions to finding such a path deformation, and so obstruct F from being the trivial element of $\pi_0 \mathcal{P}$.

2.2.1 Generic paths of functions in \mathcal{F}

We recall genericity theorems of Cerf and Hatcher-Wagoner for paths $f_t \in \mathcal{F}$. Hatcher and Wagoner also consider 2-parameter families of functions in \mathcal{F} ; this is important to show that the various invariants are well defined however we will not discuss them here and refer the reader to [HW73].

Following [Cer70], a generic path $f_t \in \mathcal{F}$ has the following properties. Except for at finitely many discrete values of t , f_t is a Morse function with no j/i trajectories for $j \leq i$. At the exceptional values of t , f_t may additionally have either a single birth-death critical value (corresponding to the creation or cancellation of an $i + 1$, i -handle pair), or a single i/i trajectory (corresponding to a handle slide); otherwise f_t has only Morse critical points, and no other j/i trajectories for $j \leq i$ as above.

We can display the critical value information of f_t as follows.

Definition 2.2.1. The *Cerf graphic* of a generic path f_t is the subset of $I \times I$ given by

$$\bigcup_{t \in I} t \times \{\text{critical values of } f_t\} \subset I \times I.$$

We further annotate our Cerf graphics with an arrow whenever there is an i/i trajectory; we draw this arrow between the two critical values. See Figure 2.5 for an example.




Figure 2.5: A Cerf graphic annotated by arrows to show trajectories between handles of the same index. This path satisfies the 1-parameter ordering condition of Proposition 2.2.2. Left to right we see a birth, an i/i handle slide, an $i + 1$ crossing, and i crossing, then an $i + 1/i + 1$ handle slide.

Just as we can deform Morse functions to be self-indexing, we can deform paths of functions to have desirable properties. We recall the one-parameter ordering theorem of Cerf [Cer70]; see [HW73] for a detailed treatment of this.

Proposition 2.2.2 (One-parameter ordering). *Let $f_t \in \mathcal{F}$ be a path whose endpoints f_0 and f_1 are Morse functions with ordered, distinct critical values (i.e. almost self-indexing, but perturbed so the critical values are distinct). We can deform this family fixing the endpoints so that f_t is a Morse function with ordered, distinct critical values for all but finitely many values of t . At these exceptional values of t , either two critical points of index i have the same critical value (shown as a crossing on the Cerf Graphic), or there is a single birth-death point, or a single i/i trajectory. Note that this necessarily means that each $(i + 1)/i$ birth-death critical point has critical value between the critical values*

of index i and those of index $i + 1$. Further, we may arrange that the birth and death points are independent, meaning that there are no trajectories between any birth/death point and another critical value. If a path has these properties then we say it satisfies the one-parameter ordering condition. See Figure 2.5 for the Cerf graphic of a 1-parameter family satisfying the one-parameter ordering condition.

If the endpoints of f_t , f_0 and f_1 contain only index i and $i + 1$ critical points for some i , then in fact we can do better

Theorem 2.2.3. [HW73, Theorem 3.1] *Suppose $f_t \in \mathcal{F}$ is a path such that f_0 and f_1 are Morse functions with only index i and $i + 1$ critical points for some $2 \leq i \leq n - 2$. Then we may deform f_t fixing the boundary so that for all $t \in I$, f_t has only critical points of index i and $i + 1$. If additionally f_0 and f_1 are Morse functions with ordered, distinct critical values then we can further deform f_t fixing the endpoints so that it additionally satisfies the 1-parameter ordering condition.*

In particular when $f_0, f_1 \in \mathcal{E}$ we can deform f_t to satisfy the conditions of Theorem 2.2.3 for any i of our choosing, provided $2 \leq i \leq n - 2$.

Remark 2.2.4. If additionally $f_0, f_1 \in \mathcal{E}$, then since the birth and death points can be made independent by Theorem 2.2.2, we may arrange that the births occur in a neighbourhood of 0, and that the deaths occur in a neighbourhood of 1, and that in this neighbourhood no handle slides, critical value crossings, or handle slides occur; see Figure 2.6. For details on how to do this, see [HW73, Chapter 1, §7].




Figure 2.6: As in Remark 2.2.4 any path with $f_0, f_1 \in \mathcal{E}$ can be deformed to satisfy the 1-parameter ordering condition, to have only critical points of index i and $i + 1$ for some i , and to have the pictured Cerf graphic in neighbourhoods of 0 and 1. Here we depict 3 births and 3 deaths, but an arbitrary number is possible (note that the number of births and deaths must be the same).

2.3 Geometric picture and conventions in dimension 4

In this section we describe the geometric picture for paths of functions $f_t \in \mathcal{F}$ in dimension 4 and set some conventions. We will henceforth assume that $f_0, f_1 \in \mathcal{E}$, i.e. they have no critical points.

By Theorem 2.2.3, after a deformation, we may assume f_t only has index 2 and 3 critical points, and that the critical values are ordered, indeed we assume that $f_t(p) < 1/2$ for critical points of index 2, $f_t(p) > 1/2$ for critical points of index 3, and $f_t(p) = 1/2$ for births and deaths. We may also assume all births happen before time ε and all deaths happen after time $1 - \varepsilon$ as in Remark 2.2.4, for some $\varepsilon \in (0, 1/4)$.

Away from births, deaths, and handle slides, f_t gives a handle decomposition of $X \times I$, relative to $X \times 0$, with only 2-handles and 3-handles. We wish to look at the “middle 4-manifold” after attaching the 2-handles, but before attaching the 3-handles. We make the following identification

$$\bigcup_{t \in [\varepsilon, 1-\varepsilon]} f_t^{-1}(1/2) = V \times [\varepsilon, 1 - \varepsilon]$$

where $V \times t = f_t^{-1}(1/2) \cong X \#^m (S^2 \times S^2) \forall t$. Here m is the number of births (also deaths), and we see one copy of $S^2 \times S^2$ for each 2-3 handle pair. Directly after the births we see the belt sphere for a 2-handle as $S^2 \times p$ and the attaching spheres of the corresponding 3-handle as $q \times S^2$ in each $S^2 \times S^2$ summand. As we move forward in the t direction we see an isotopy of these attaching spheres. The spheres will also change when a handle slide occurs. To keep track of these spheres we establish the following conventions.

1. We denote the 2-handle belt spheres in any given t slice by $A_1^t, \dots, A_n^t \subset V \times t$. We refer to these collectively as the A -spheres.
2. We denote the 3-handle attaching spheres in any given t slice by B_1^t, \dots, B_n^t . We refer to these collectively as the B -spheres.
3. We orient the A_i^t s and B_i^t s so that the intersection between A_i^t and B_i^t directly after their birth is positive. There is a consistent choice of orientation for all t so that the 3-manifolds $\cup_t A_i^t$ and $\cup_t B_i^t$ are oriented; note that these 3-manifolds have S^2 boundary components at the handle slides.

4. We pick a basepoint $* \in V$; this gives us a basepoint for $V \times I$ or indeed $X \times I \times I$ by taking $* \times 0 \in V \times I \subset X \times I \times I$. We will often abusively refer to $*$ as the base point in any given t slice; really we mean $* \times t \in V \times t$. If we refer to a path to this basepoint, we implicitly take the further path to the “true basepoint” $* \times 0$ by taking a path through $* \times I$.
5. We make a continuous choice of basepoint in the A -spheres and B -spheres for each value of t ; that is $*_{A_i^t} \in A_i^t \subset V \times t$, $*_{B_i^t} \in B_i^t \subset V \times t$.
6. We make a continuous choice for all t of path from the basepoint of V to the basepoints of the spheres; $\alpha_i^t \subset V \times t$ from $*$ to $*_{A_i^t} \in A_i^t$ and $\beta_i^t \subset V \times t$ from $*$ to $*_{B_i^t} \in B_i^t$. We do so such that directly after the creation of each pair, we have $\alpha_i^t \cdot \beta_i^{t-1} = 1 \in \pi_1 V$.

2.3.1 Handle slides

At each 2/2 trajectory we see a handle slide between the 2-handles, and similarly at each 3/3 trajectory we see a handle slide between the 3-handles; we may assume that there are finitely many of these and that they occur at distinct values of t in $(\varepsilon, 1 - \varepsilon)$.

We describe the effect of the handle slides on the A -spheres and B -spheres. Consider a 3/3 handle slide at time t , where we slide the handle attached to $B_k^{t-\delta}$ over the handle attached to $B_j^{t-\delta}$. After the handle slide $B_j^{t+\delta} = B_j^{t-\delta}$, while $B_k^{t+\delta}$ is a connect sum of $B_k^{t-\delta}$ to a parallel copy of $B_j^{t-\delta}$.

The picture at 2/2 handle slides is similar. In this case we slide the handle with belt sphere $A_k^{t-\delta}$ over the handle with belt sphere $A_j^{t-\delta}$ at time t . Then $A_k^{t+\delta} = A_k^{t-\delta}$, while $A_j^{t+\delta}$ is a connect sum of $A_j^{t-\delta}$ to a parallel copy of $A_k^{t-\delta}$. We can see this by turning the handle decomposition upside down and considering the belt spheres of the 2-handles as attaching regions of some 3-handles; note that after turning upside down the A_j 3-handle is being slid over the A_k 3-handle.

Remark 2.3.1. In fact, the handle slides are determined by the (framed) arc in $V \times (t - \delta)$ from $A_k^{t-\delta}$ to $A_j^{t-\delta}$ or $B_k^{t-\delta}$ to $B_j^{t-\delta}$. We refer to this choice of arc as the *handle slide arc*.

2.3.2 Intersections of spheres

Throughout f_t , by our genericity assumptions the A_i s are disjoint from each other, as are the B_j s. There may be intersections between the A_i s and the B_j s however. Again by genericity of f_t , we can also assume that the A_i s and B_j s intersect transversely.

Initially after the births, the handles A_i^t and B_j^t intersect in $\delta_{i,j}$ points. At later t this may no longer be true. Note that the intersections form a 1-manifold in $V \times I$, and that the endpoints of any arcs in the 1-manifold occur at the births, the deaths or the handle slides.

Away from the handle slides we see a regular homotopy of the spheres which restricts to an ambient isotopy of each of the families $\{A_i^t\}$ and $\{B_i^t\}$. Hence new intersections are introduced and removed by finger moves and Whitney moves between some A sphere and some B sphere.

Remark 2.3.2. When we wish to construct paths of functions $f_t \in \mathcal{F}$, we can do so by creating a 1-parameter family of handle structures. The rules for constructing 1-parameter families of handle structures are the same as those for creating generic 1-parameter families of functions in \mathcal{F} ; we can create cancelling pairs of i , $(i+1)$ -handles, perform handle slides, and perform isotopy of the attaching regions of handles, and cancel pairs of handles that intersect in a single point. Given such a 1-parameter family of handle structures there is certainly some 1-parameter family in $f_t \in \mathcal{F}$ which induces this family of handle structures. In dimension 4, when we only have 2 and 3 handles we can do all of this by considering deformation of the A and B spheres in the middle level; see [Qui86] for Quinn's treatment of this in dimension 4.

2.4 Review of Hatcher and Wagoner's $\text{Wh}_2(\pi_1 X)$ invariant Σ

In this section we recall the definition of the map $\Sigma: \pi_0 \mathcal{P} \rightarrow \text{Wh}_2(\pi_1 X)$ of Hatcher and Wagoner, and the key result which we will use, namely the reduction to eyes for elements in the kernel of Σ ; see [HW73, Chapter VI].

2.4.1 Algebra of Wh_2

We begin by recalling some algebraic definitions.

Let $\Lambda = \mathbb{Z}[\pi_1 X]$. Let $\text{GL}(\Lambda) = \lim_{n \rightarrow \infty} \text{GL}_n(\Lambda)$. For $\lambda \in \Lambda$ let $e_{i,j}^\lambda \in \text{GL}(\Lambda)$ be the matrix which is the identity on the diagonal, has λ in the (i, j) position, and is zero elsewhere. We call $e_{i,j}^\lambda$ an *elementary matrix*, and let $\text{E}(\Lambda) \subset \text{GL}(\Lambda)$ be the subgroup generated by the elementary matrices.

One can easily verify the following relations in $\text{E}(\Lambda)$:

- (i) $e_{i,j}^\lambda \cdot e_{i,j}^\mu = e_{i,j}^{\lambda+\mu}$,
- (ii) $[e_{i,j}^\lambda, e_{k,l}^\mu] = 0$ for $i \neq l$ and $j \neq k$, and
- (iii) $[e_{i,j}^\lambda, e_{j,l}^\mu] = e_{i,l}^{\lambda\mu}$ for i, j, l distinct.

This motivates the following definition.

Definition 2.4.1. The *Steinberg group* $\text{St}(\Lambda)$ is the group freely generated by symbols $x_{i,j}^\lambda$ for $i, j \in \mathbb{N}$ and $\lambda \in \Lambda$ subject to the relations

- (i) $x_{i,j}^\lambda \cdot x_{i,j}^\mu = x_{i,j}^{\lambda+\mu}$,
- (ii) $[x_{i,j}^\lambda, x_{k,l}^\mu] = 0$ for $i \neq l$ and $j \neq k$, and
- (iii) $[x_{i,j}^\lambda, x_{j,l}^\mu] = x_{i,l}^{\lambda\mu}$ for i, j, l distinct.

Note that we have a surjective homomorphism $\pi: \text{St}(\Lambda) \rightarrow \text{E}(\Lambda)$ sending $x_{i,j}^\lambda \mapsto e_{i,j}^\lambda$.

We define $K_2(\Lambda)$ to be the kernel of $\pi: \text{St}(\Lambda) \rightarrow \text{E}(\Lambda)$. Hence we have the short exact sequence:

$$0 \rightarrow K_2(\Lambda) \rightarrow \text{St}(\Lambda) \xrightarrow{\pi} \text{E}(\Lambda) \rightarrow 0.$$

For $g \in \pi_1 X$ let $w_{i,j}^{\pm g} = x_{i,j}^{\pm g} x_{j,i}^{\mp g^{-1}} x_{i,j}^{\pm g}$, and let $W(\pm\pi_1 X) \subset \text{E}(\Lambda)$ be the subgroup generated by the words $w_{i,j}^{\pm g}$. Then we define the *second Whitehead group* to be,

$$\text{Wh}_2(\pi_1 X) = K_2(\Lambda) \text{ mod } W(\pm\pi_1 X) \cap K_2(\Lambda)$$

In order to define Σ we need the following lemma.

Lemma 2.4.2 ([HW73, Chapter III, Lemma 1.6]). *Let $P \in \text{GL}(\Lambda)$ be a permutation matrix, and let $D \in \text{GL}(\Lambda)$ be diagonal with entries in $\pm\pi_1 X$. Then there exists some $w \in W(\pm\pi_1 X)$ such that $\pi(w) = P \cdot D$.*

This follows from the fact that $P \cdot D$ can be written as a product

$$P \cdot D = \prod_k e_{i_k, j_k}^{\pm g_k} e_{j_k, i_k}^{\mp g_k^{-1}} e_{i_k, j_k}^{\pm g_k}.$$

2.4.2 Definition of Σ

We now recall the definition of $\Sigma: \mathcal{P} \rightarrow \text{Wh}_2(\pi_1 X)$. For full details see [HW73, Chapter IV], and see [HW73, Chapter V, §6] for the version of the definition we give here. The approach is to construct a map (which we also abusively call Σ)

$$\Sigma: \pi_1(\mathcal{F}, \mathcal{E}) \rightarrow \text{Wh}_2(\pi_1 X),$$

and use Π to identify $\pi_0 \mathcal{P}$ with $\pi_1(\mathcal{F}, \mathcal{E})$.

As in Remark 2.2.4 we deform f_t so that it satisfies 1-parameter ordering, has only critical points of index i and $i + 1$ for some $2 \leq i \leq n - 2$, that the birth points appear in time interval $(0, \varepsilon)$, that all death points appear in the time interval $(1 - \varepsilon, 1)$ and that no critical lines cross and no handle slides occur in either of these intervals, so that in these intervals the Cerf graphic is as in Figure 2.6.

We fix a base point $*$ in $X \times I$ and a local orientation at $*$. For each birth point $c \in X \times I$ we choose a label $j \in \mathbb{N}$ for this critical point, and a path in $\gamma_c \subset X \times I$ from $*$ to c . For $t \in [\varepsilon, 1 - \varepsilon]$ denote the index $i + 1$ and i critical points created by this birth by z_t^j and b_t^j respectively. Using the path γ_c we can obtain a continuous choice of path from $*$ to z_t^j and b_t^j for each t , denoted $\gamma_{b_t^j}$ and $\gamma_{z_t^j}$ respectively. We also choose an orientation of the stable sets. This gives a basis for $C_i(f_t, \eta_t)$ as in Section 2.1.1 and a matrix $\partial_{i+1}^t = \partial_t \text{GL}(\mathbb{Z}[\pi_1 X])$.

After the birth points ∂_ε is the identity matrix, as there is precisely one trajectory from z^j to b^j for each j .

The matrix ∂_t remains constant until passing an i/i or $i + 1/i + 1$ trajectory. On passing an $i + 1/i + 1$ trajectory φ from z_t^j to z_k^t , with associated $\pm \pi_1 X$ element given by $\lambda = \gamma_{z_t^j}^{-1} \cdot \varphi \cdot \gamma_{z_k^t}^{-1} \in \pm \pi_1 X$ (with sign depending on the orientations), we have that

$$\partial_{t+\delta} = \partial_{t-\delta} \circ e_{j,k}^\lambda.$$

On passing an i/i trajectory φ from b_j^t to b_k^t with associated $\pm \pi_1 X$ element given by $\lambda =$

$\pm \gamma_{b_t^j} \circ \varphi \circ \gamma_{b_t^k}^{-1} \in \pm \pi_1 X$, we have

$$\partial_{t+\delta} = \left(e_{j,k}^\lambda \right)^{-1} \circ \partial_{t-\delta} = e_{j,k}^{-\lambda} \circ \partial_{t-\delta}$$

Hence just before the death points at time $1 - \varepsilon$ we have,

$$\begin{aligned} \partial_{1-\varepsilon} &= e_{j'_m, k'_m}^{-\lambda'_m} \cdots e_{j'_2, k'_2}^{-\lambda'_2} e_{j'_1, k'_1}^{-\lambda'_1} \partial_\varepsilon e_{j_1, k_1}^{\lambda_1} e_{j_2, k_2}^{\lambda_2} \cdots e_{j_l, k_l}^{\lambda_l} \\ &= e_{j'_m, k'_m}^{-\lambda'_m} \cdots e_{j'_2, k'_2}^{-\lambda'_2} e_{j'_1, k'_1}^{-\lambda'_1} e_{j_1, k_1}^{\lambda_1} e_{j_2, k_2}^{\lambda_2} \cdots e_{j_l, k_l}^{\lambda_l}. \end{aligned}$$

On the other hand, at time $1 - \varepsilon$ each critical point b_j will cancel with some critical point z_j , so we know there is precisely one trajectory between them along some arc $\lambda \in \pm \pi_1 X$, and no other trajectories. Hence,

$$\partial_{1-\varepsilon}(z_j) = \lambda \cdot b_k$$

for some k . Hence we can write $\partial_{1-\varepsilon}$ as a product,

$$\partial_{1-\varepsilon} = P \cdot D$$

where P is a permutation matrix, and D is diagonal with entries in $\pm \pi_1 X$. By Lemma 2.4.2, there exists $w \in W(\pm \pi_1 X)$ such that $\pi(w) = P \cdot D$.

We can now define $\Sigma: \pi_1(\mathcal{F}, \mathcal{E}) \rightarrow \text{Wh}_2(\pi_1 X)$ by sending the path f_t to

$$x_{j'_m, k'_m}^{-\lambda'_m} \cdots x_{j'_2, k'_2}^{-\lambda'_2} x_{j'_1, k'_1}^{-\lambda'_1} x_{j_1, k_1}^{\lambda_1} x_{j_2, k_2}^{\lambda_2} \cdots x_{j_l, k_l}^{\lambda_l} \cdot w^{-1} \text{ mod } W(\pm \pi_1 X) \cap K_2(\Lambda) \in \text{Wh}_2(\pi_1 X).$$

Note that $\pi \left(x_{j'_m, k'_m}^{-\lambda'_m} \cdots x_{j'_1, k'_1}^{-\lambda'_1} x_{j_1, k_1}^{\lambda_1} \cdots x_{j_l, k_l}^{\lambda_l} \cdot w^{-1} \right) = \partial_{1-\varepsilon} \partial_{1-\varepsilon}^{-1} = 1$ so this is indeed an element of $K_2(\Lambda)$.

Hatcher and Wagoner that in dimension $n \geq 4$ that $\Sigma(f_t)$ does not depend on the various choices made above, and indeed is independent of the homotopy class of the path in $(\mathcal{F}, \mathcal{E})$ as required. See [HW73, Chapter IV, §3] for the proof that Σ is well defined.

Remark 2.4.3. If $P \cdot D$ were the identity matrix, then the product

$$x_{j'_m, k'_m}^{-\lambda'_m} \cdots x_{j'_1, k'_1}^{-\lambda'_1} x_{j_1, k_1}^{\lambda_1} \cdots x_{j_l, k_l}^{\lambda_l}$$

would give an element of $K_2(\pi_1 X)$. In the case that it is not, we could deform the 1-parameter family so that $P \cdot D$ is the identity by adding a trivial pseudo-isotopy corresponding to word the w defined above. Indeed by [HW73, Chapter IV, Lemma 2.7] for

$n \geq 4$, for every such $w \in W(\pm\pi_1 X)$ there exists a path f_t with handle slides corresponding to w , such that f_t is isotopic to the trivial path fixing the end points.

2.4.3 Reduction to eyes and $\ker \Sigma$

Hatcher and Wagoner prove the following about paths in the kernel of Σ .

Theorem 2.4.4 ([HW73, Chapter VI, Theorem 2]). *Let $n \geq 4$. Given a 1-parameter family f_t representing a class in $\pi_1(\mathcal{F}, \mathcal{E})$ for which $\Sigma(f_t) = 0$, there is a deformation of f_t fixing the end points to a 1-parameter family f'_t whose Cerf graphic consists of a nested collection of eyes, with only index i and $i + 1$ critical points for some $2 \leq i \leq n - 2$, and no i/i or $i + 1/i + 1$ trajectories; this necessarily means that each $(i + 1)$ -handle cancels at a death point with the i -handle it was created with. See Figure 2.7. We can also assume that the births and deaths remain independent.*

Hatcher and Wagoner in fact prove that when $n \geq 5$ any one-parameter family f_t with $\Sigma(f_t) = 0$ can be deformed to a 1-parameter family consisting of only a *single* eye. As noted by Quinn in [Qui86] the reduction to multiple eyes holds in 4-dimensions, and only the step reducing from many eyes to a single eye requires dimension $n \geq 5$.




Figure 2.7: A Cerf graphic that is a collection of nested eyes.

2.5 Stable surjectivity of Σ in dimension 4

In [HW73] they show the following result.

Theorem 2.5.1 ([HW73, Theorem 2, Chapter VI]). $\Sigma: \pi_0 \mathcal{P} \rightarrow \text{Wh}_2(\pi_1 X)$ is surjective in dimension $n \geq 5$.

The proof is simple, we recall it here to motivate our upcoming proof.

Proof. Given $x \in \text{Wh}_2(\pi_1 X)$, let $x_{j_1, k_1}^{s_1 \lambda_1} \dots x_{j_m, k_m}^{s_m \lambda_m} \in K_2(\mathbb{Z}[\pi_1 X])$ with $\lambda_l \in \pi_1 X$ and $s_l \in \{+1, -1\}$ be a word representing x ; note that any word in $K_2(\mathbb{Z}[\pi_1 X])$ can be written in this way.

We build a path f_t in \mathcal{F} as in Remark 2.3.2. Let $N = \max(\max_l(j_l), \max_l(k_l))$. Starting with the trivial element $p \in \mathcal{F}$, we first create N pairs of cancelling i and $(i+1)$ -handles for some $2 \leq i \leq n-2$, and label them $1, \dots, N$.

We then perform m handle slides of the $(i+1)$ -handles over $(i+1)$ -handles. For $l = 1, \dots, m$ we slide the j_l th $(i+1)$ -handle over the k_l th $(i+1)$ -handle along an arc so that the associated element of $\pi_1 X$ is λ_l . For each handle slide we can also choose to perform an oriented or unoriented handle slide when $s_l = +1$ or $s_l = -1$ respectively. After these handle slides, the resulting differential is

$$\partial_{i+1} = e_{j_1, k_1}^{s_1 \lambda_1} \dots e_{j_m, k_m}^{s_m \lambda_m} = \pi(x_{j_1, k_1}^{s_1 \lambda_1} \dots x_{j_m, k_m}^{s_m \lambda_m}) = 1.$$

Now as in the proof of the s -cobordism theorem we may realise this differential geometrically; that is in the middle level, we know that algebraically the attaching region for the j th $(i+1)$ -handle and belt sphere for the k th i -handle intersect in $\delta_{j,k}$ points. We make use of the Whitney trick to deform these spheres (and correspondingly the handle decomposition) so that they intersect in $\delta_{j,k}$ points geometrically. Note that the Whitney trick requires dimension $n \geq 5$.

Now the handles intersect geometrically in $\delta_{j,k}$ points. We may cancel the pairs of handles, creating N deaths. After the deaths, the resulting Morse function is an element of \mathcal{E} as it has no critical points. Hence we have described a path in $\pi_1(\mathcal{F}, \mathcal{E})$ as required. It is clear that this path has $\Sigma(f_t) = x_{j_1, k_1}^{s_1 \lambda_1} \dots x_{j_m, k_m}^{s_m \lambda_m}$. \square

This proof does not work in dimension 4 since we cannot apply the Whitney trick. In this section we present an proof of a stable version of this theorem in dimension 4.

We first prove a further useful lemma for constructing transverse spheres.

Lemma 2.5.2. *Let X be a 4-manifold. Consider a handle decomposition of $X \times I$ relative to $X \times 0$ with N 3-handles. Denote the attaching regions for the 3-handles by $B_1, \dots, B_N \subset V$ where V is the middle level between the 2 and 3 handles. Suppose there are disjointly embedded spheres $T_1, \dots, T_N \subset V$ such that B_i intersects T_j in $\delta_{i,j}$ points.*

Assume that T_1, \dots, T_N have trivial normal bundle. Suppose we now perform some number of handle slides between the 3-handles to obtain a new set of 3-handle spheres B'_1, \dots, B'_N . We claim there exists disjointly embedded spheres $T'_1, \dots, T'_N \subset V$ such that B'_i intersects T'_j in $\delta_{i,j}$ points, and that T'_i and T'_j are disjoint for all i and j . We also claim that we may choose T'_1, \dots, T'_N to have trivial normal bundle.

Proof. We prove the claim by induction on the number of handle slides. Denote the 3-handle attaching spheres after k handle slides by B_1^k, \dots, B_N^k , we construct T_1^k, \dots, T_N^k which are dual to B_1^k, \dots, B_N^k and have trivial normal bundle. Since T_i has trivial normal bundle, letting T_i^0 be a parallel copy of T_i is sufficient for $k = 0$.




Figure 2.8: Performing the $(k + 1)$ th handle slide, then performing the Norman trick to remove the extra intersection.

Suppose we have a family of spheres T_1^k, \dots, T_N^k with trivial normal bundle, which are dual to B_1^k, \dots, B_N^k and disjoint from T_1, \dots, T_N . We now perform the $(k + 1)$ th handle slide, which we may assume slides B_i^k over B_j^k using some arc $\gamma \subset V$ from B_i^k to B_j^k . By general position we may perform an isotopy of γ so that it is disjoint from T_1^k, \dots, T_N^k and T_1, \dots, T_N (it is necessarily disjoint from B_1^k, \dots, B_N^k except at the endpoints). Now $B_i^{k+1} = B_i^k \#_\gamma B_j^k$. Hence B_i^{k+1} intersects T_i^k in a single point p , and T_j^k in a single point q . We remove the intersection between B_i^{k+1} and T_j^k by choosing an arc in B_i^{k+1} from q to p , and tubing T_j^k to a parallel copy of T_i^k using the Norman trick. We call the resulting sphere T_j^{k+1} ; see Figure 2.8. Note that B_i^{k+1} may intersect T_1, \dots, T_N in finitely many points so we must ensure we choose the arc from q to p so that it misses these intersections, and so we do not

introduce any intersections with T_1, \dots, T_N . Since T_i^k is disjoint from T_1, \dots, T_N provided we choose a small enough neighbourhood in which to take the parallel copy we can ensure that the T_j^{k+1} sphere does not intersect T_1, \dots, T_N . Taking $T_r^{k+1} = T_r^k$ for $r \neq j$ we see that $T_1^{k+1}, \dots, T_N^{k+1}$ are a dual family for $B_1^{k+1}, \dots, B_N^{k+1}$ as required, and do not intersect T_1, \dots, T_N .

We also note that T_1^k, \dots, T_N^k have trivial normal bundle, since tubing spheres with trivial normal bundle together results in a sphere with trivial normal bundle. \square

Remark 2.5.3. By turning the handle decompositions in the above proof upside down, we can prove the same fact about the belt spheres of 2-handles in V , performing 2-handle slides instead of 3-handle slides.

2.5.1 Proof of stable surjectivity in dimension 4

In this section we prove that Σ is stably surjective in dimension 4.

Theorem E. *Let X be a compact 4-manifold and $x \in \text{Wh}_2(\pi_1 X)$. There exists N , and a pseudo-isotopy F of $X \#^N S^2 \times S^2$ such that*

$$\Sigma(F) = x \in \text{Wh}_2(\pi_1(X \#^N S^2 \times S^2)) = \text{Wh}_2(\pi_1 X).$$

Proof. As in the proof of Theorem 2.5.1 we consider the word $x_{i_1, j_1}^{s_1 \lambda_1} \dots x_{i_m, j_m}^{s_m \lambda_m} \in K_2(\mathbb{Z}[\pi_1 X])$ representing x , where $s_k \in \{+1, -1\}$, $\lambda_k \in \pi_1(X)$, and $i_k, j_k \in 1, \dots, N$ for some N .

Let $Y_N = \#^N S^2 \times S^2$ and let $X' = X \# Y_N$. We claim that N is sufficient to find a pseudo-isotopy of X' such that

$$\Sigma(F) = \left[x_{i_1, j_1}^{s_1 \lambda_1} \dots x_{i_m, j_m}^{s_m \lambda_m} \right] = x \in \text{Wh}_2(\pi_1(X \#^N S^2 \times S^2)).$$

To show the existence of such a pseudo-isotopy we construct a path in $\mathcal{F}(X')$ by considering a deformation of handle structures of $X' \times I$ as in Remark 2.3.2. We start with the trivial element of $\mathcal{F}(X')$ and correspondingly the trivial handle structure for $X' \times I$.

We denote the X summand of $X' = X \# Y_N$ by $\widehat{X} = X \setminus B^3 \subset X'$. We initially describe a handle deformation just in $\widehat{X} \times I$ and leave the handle structure of

$$(Y_N \setminus B^3) \times I \subset X' \times I$$

fixed. Note that $\pi_1(\widehat{X}) = \pi_1(X)$.

We first create N cancelling 2 and 3-handle pairs in \widehat{X} and then perform m handle slides, again within \widehat{X} , of 3-handles over 3-handles in accordance with the word

$$x_{i_1, j_1}^{s_1 \lambda_1} \cdots x_{i_m, j_m}^{s_m \lambda_m} \in K_2(\mathbb{Z}[\pi_1 X]) = K_2(\mathbb{Z}[\pi_1 X']).$$

As usual we consider the middle level $V = X' \#^N S^2 \times S^2$ between the 2-handles and 3-handle. In order to differentiate this copy of $\#^N S^2 \times S^2$ from Y_N , we let $Z_N = \#^N S^2 \times S^2$ and let $V = X' \# Z_N$. Since the cancelling pairs were created in $\widehat{X} \times I$, we consider the middle level of $\widehat{X} \times I$, namely $\widehat{V} = \widehat{X} \# Z_N \subset V$. Note that

$$V = X' \# Z_N = X \# Y_N \# Z_N = (\widehat{X} \# Z_N) \cup (Y_N \setminus B^4) = \widehat{V} \cup (Y_N \setminus B^4)$$

Note also that the 2-handle belt spheres and 3-handle attaching spheres lie within \widehat{V} throughout the handle slides.

We fix a time t_0 after the handle slides have taken place and work in this fixed time t_0 . In the middle level at time t_0 let $A_1, \dots, A_N \subset \widehat{V}$ be the belt spheres of the 2-handles and $B_1, \dots, B_N \subset \widehat{V}$ be the attaching spheres for the 3-handles. Note that since the 2-handle spheres did not move during the handle slides, they are still $S^2 \times p$ slices of $Z_N = \#^N S^2 \times S^2$.

Since prior to the handle slides the i th 3-handle attaching sphere and the j th 2-handle belt sphere (which since it remained fixed is $A_j \subset V$) intersect in $\delta_{i,j}$ points, by Lemma 2.5.2 there exist spheres C_1, \dots, C_N with trivial normal bundle such that B_i and C_j intersect in $\delta_{i,j}$ points, and that A_i and C_j are disjoint for all i, j .

As $x_{i_1, j_1}^{s_1 \lambda_1} \cdots x_{i_m, j_m}^{s_m \lambda_m} \in K_2(\mathbb{Z}[\pi_1 X])$, the differential ∂_3 is the identity, so A_i and B_j have algebraic intersection $\delta_{i,j} \in \mathbb{Z}[\pi_1 X]$. This means that we can pair all of the intersections up (except for a single intersection for each A_i, B_i pair) so that for each of the paired intersection points $p, q \in A_i \cap B_j$ the intersection as measured in $\pm \pi_1 X$ is $+\gamma$ for p and $-\gamma$ for q for some $\gamma \in \pi_1 X$. We can also pick Whitney arcs for each pair; note that removing disjointly embedded arcs from a sphere does not make the sphere disconnected, so we can pick disjoint Whitney arcs for all of the pairings. The resulting Whitney circles for each pair of intersections vanish in $\pi_1 X = \pi_1 \widehat{V}$, and so there exists a Whitney disc $W \subset \widehat{V}$ for each pair of points. We can perform boundary twists between W and B_j so that W is a

correctly framed Whitney disc. We do this for all of the pairs, and refer to the collection of these Whitney discs as the W -discs.

The W -discs may self intersect, and also intersect other W -discs. They may also intersect the A -spheres, the B -spheres and the C -spheres. We will manipulate these discs to construct a new family of discs which are embedded and disjoint from the A -spheres and B -spheres.

First, noting that A_i is a copy of $S^2 \times p$ in $\widehat{V} = \widehat{X} \# Z_N = \widehat{X} \#^N S^2 \times S^2$, there exists some transverse sphere for A_i , $A_i^* = q \times S^2$. For each intersection r between a W -disc and an A_i we use the Norman trick to tube the W -disc into a parallel copy of A_i^* . Doing this requires a choice of arc in A_i from r to $A_i \cap A_i^*$, when choosing this arc we ensure it misses any other intersection points in A_i and any Whitney arcs; this is possible as these do not disconnect A_i . We do this for each intersection between an A -sphere and W -disc successively. We obtain new W -discs which are disjoint from the A -spheres; note that this may introduce intersections between W -discs and W -discs, as well as new intersections between the W -discs and the B -spheres and C -spheres. The W -discs are also still correctly framed Whitney discs as the dual spheres A_i^* have trivial normal bundle.

Next, we push down the self intersections of the W -discs into the B -spheres, as in Section 2.1.2. After doing this for all intersections between the W -discs, the W -discs are disjoint and embedded and disjoint from the A -spheres, but still possibly intersect the C and B spheres. The W -discs are also still correctly framed Whitney discs.

Since everything we have done so far has been within $\widehat{V} = \widehat{X} \# Z_N$, all of the W -discs, A -spheres, B -spheres and C -spheres are disjoint from the $Y_N \setminus B^4$ subset of

$$V = (\widehat{X} \# Z_N) \cup (Y_N \setminus B^4).$$

We now make use of the $S^2 \times S^2$ s in $Y_N = \#^N S^2 \times S^2$. For each C -sphere C_j , we create a new sphere C'_j by tubing C_j into $S^2 \times p \subset S^2 \times S^2$ in Y_N . Ensuring to tube each C -sphere into a different $S^2 \times S^2$, we obtain spheres C'_1, \dots, C'_N which are disjoint and embedded. Taking $D_j = p \times S^2$ in each summand of Y_N , we also have spheres D_1, \dots, D_n such that D_i and C'_j intersect in $\delta_{i,j}$ points. Clearly the D -spheres are disjoint from the W -discs, the A -spheres and the B -spheres. Note also that the C' -spheres have trivial normal bundle, as do the D -spheres.

We can now remove intersections between the W discs and the C' -spheres by using the Norman trick; for each intersection between a W -disc and C'_j we create a new Whitney disc by tubing into a parallel copy of D_j , noting that after each tubing D_j is still disjoint from the W -discs so we may repeat the process without introducing new intersections between W -discs. After doing this for each intersection we obtain embedded W -discs which are disjoint from the A -spheres and C' -spheres but still intersect the B -spheres. The W -discs are also still correctly framed Whitney discs as each D_j has trivial normal bundle.

We now remove the intersections between the B -spheres and the W -discs; for each intersection p between a W -disc and B_j , pick an arc from p to $B_j \cap C'_j$ which is disjoint from the Whitney arcs and other intersection points (note that the Whitney arcs do not disconnect B_j so this is always possible), then use this arc to perform the Norman trick, tubing the W -disc into a parallel copy of C'_j . Note that C'_j is disjoint from the W -discs so this does not introduce new intersections between the W -discs. Since after each tubing C'_j is still disjoint from the W -discs, we may repeat the process for every intersection between B -spheres and W -discs. Note also that C'_j does not intersect the A -spheres so this does not introduce intersections between W -discs and A -spheres. At the end of this process we obtain embedded W -discs which are disjoint from all the A -spheres and B -spheres. Additionally the W -discs are still framed as C'_j has trivial normal bundle for all j .

Finally, having built embedded W -discs in V at time t_0 , we now use them to perform Whitney moves for each pair of points to remove the pairs of intersections; note that to do this we stop working in time t_0 . After performing these Whitney moves, A_i and B_j intersect in $\delta_{i,j}$ points, so we can cancel all of the handles handles. After we do this the resulting handle structure has no handles, so the corresponding Morse function lies in $\mathcal{E}(X \#^N S^2 \times S^2)$. Hence we have built a path in $\mathcal{F}(X \#^N S^2 \times S^2)$ with endpoints in $\mathcal{E}(X \#^N S^2 \times S^2)$. It is clear that $\Sigma(f_t) = x_{i_1, j_1}^{s_1 \lambda_1} \cdots x_{i_m, j_m}^{s_m \lambda_m}$ as required. \square

2.6 The $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$ invariant Θ

In Section 2.4 we described the map Σ , which in dimension $n \geq 4$ is a complete obstruction to reducing a 1-parameter family f_t with endpoints in \mathcal{E} to one whose Cerf graphic is a collection of nested eyes as in Figure 2.7. That is, any pseudo-isotopy in $\ker \Sigma$ is represented by such a 1-parameter family.

In this section we recall the definition of the map described by Hatcher and Wagoner in [HW73, Chapter VII]

$$\Theta: \ker \Sigma \rightarrow \text{Wh}_1(\pi_1(X); \mathbb{Z}_2 \times \pi_2(X)) / \chi(K_3 \pi_1 X).$$

In dimension ≥ 6 , Hatcher and Wagoner (with later clarification by Igusa in [Igu84]) show that Θ is a complete obstruction to removing all the eyes from such a nested collection of eyes (at which point we are left with an empty Cerf graphic, and so a path in \mathcal{E}). Hence, in dimension ≥ 6 , Σ together with Θ provide a complete obstruction to a pseudo-isotopy being isotopic to an isotopy.

When $k_1 X = 0$, Igusa in [Igu84] describes a map

$$\Theta_\sigma: \pi_0 \mathcal{P} \rightarrow \text{Wh}_1(\pi_1(X); \mathbb{Z}_2 \times \pi_2(X))$$

dependent on a choice of section $\sigma: X_{(1)} \rightarrow X_{(2)}$. We will not address this extension here. When restricted to $\ker \Sigma$, Igusa's map $\Theta_\sigma|_{\ker \Sigma}$ agrees with the Θ described in [HW73, Chapter VII] and is independent of the section σ .

2.6.1 Definition of $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$

Let Γ be an abelian group acted on by a group π . For our purposes Γ will be the group $\mathbb{Z}_2 \times \pi_2(X)$ and π will be $\pi_1(X)$, with the usual action on $\pi_2(X)$ and the trivial action on \mathbb{Z}_2 . We write elements of $R = \Gamma[\pi] \times \mathbb{Z}[\pi]$ as finite formal sums $\sum_i (\alpha_i + n_i) \sigma_i$ for $\alpha_i \in \Gamma$, $\sigma_i \in \pi$, $n_i \in \mathbb{Z}$. We give R a ring structure via

$$\left(\sum_i (\alpha_i + n_i) \sigma_i \right) \cdot \left(\sum_j (\beta_j + m_j) \tau_j \right) = \sum_{i,j} (n_i \beta_j^{\sigma_i} + m_j \alpha_i + n_i m_j) \sigma_i \tau_j,$$

where $\alpha_i, \beta_j \in \Gamma$, $n_i, m_j \in \mathbb{Z}$, $\sigma_i, \tau_j \in \pi$ and β^σ denotes the action of σ on β . Note that $\Gamma[\pi]$ is an ideal of R . Note also that the multiplication is trivial on $\Gamma[\pi]$; that is $x \cdot y = 0$, for all $x, y \in \Gamma[\pi]$. We define

$$\text{GL}(\Gamma[\pi]) = \ker (\text{GL}(R) \rightarrow \text{GL}(R/\Gamma[\pi])) = \{I + A \mid A \text{ has entries in } \Gamma[\pi]\}$$

Note that R and $R/\Gamma[\pi]$ are rings, and here $\text{GL}(R)$ and $\text{GL}(R/\Gamma[\pi])$ denote the usual general linear group of matrices in these rings (that is, the union over n of $\text{GL}_n(R)$ and $\text{GL}_n(R/\Gamma[\pi])$ respectively). Note also that $I + A$ is always in $\text{GL}(R)$ if A has entries in

$\Gamma[\pi]$, since

$$(I + A) \cdot (I - A) = I - A \cdot A = I,$$

noting that $A \cdot A = 0$ as multiplication in the ideal $\Gamma[\pi]$ is trivial.

Definition 2.6.1. We define $K_1\Gamma[\pi]$ by $K_1\Gamma[\pi] = \text{GL}(\Gamma[\pi])/[\text{GL}(R), \text{GL}(\Gamma[\pi])]$

Hatcher proves the following.

Proposition 2.6.2 ([Hat73, Proposition 1.1]). *The trace map*

$$\text{tr} : K_1\Gamma[\pi] \rightarrow \Gamma[\pi]/\langle ar - ra \mid a \in \Gamma[\pi], r \in R \rangle$$

is an isomorphism. The subgroup $\langle ar - ra \mid a \in \Gamma[\pi], r \in R \rangle$ can also be expressed as $\langle \alpha\sigma - \alpha^\tau\tau\sigma\tau^{-1} \mid \alpha \in \Gamma, \tau, \sigma \in \pi \rangle$, where α^τ denotes the result of acting on α by τ . Hence we obtain

$$K_1\Gamma[\pi] \cong \Gamma[\pi]/\langle \alpha\sigma - \alpha^\tau\tau\sigma\tau^{-1} \mid \alpha \in \Gamma, \tau, \sigma \in \pi \rangle.$$

We can now define the first Whitehead group of the pair (π, Γ) .

Definition 2.6.3. $\text{Wh}_1(\pi; \Gamma) = \text{coker}(\Gamma[1] \rightarrow K_1\Gamma[\pi])$, where $\Gamma[1] \rightarrow K_1\Gamma[\pi]$ is defined by $1\alpha \mapsto [1\alpha]$

It follows easily that

Corollary 2.6.4. $\text{Wh}_1(\pi; \Gamma) = \Gamma[\pi]/\langle \alpha\sigma - \alpha^\tau\tau\sigma\tau^{-1}, \beta \cdot 1 \mid \alpha, \beta \in \Gamma, \tau, \sigma \in \pi \rangle$

Remark 2.6.5. It is clear from the definitions in this section that $\text{Wh}_1(\pi_1X; \mathbb{Z}_2 \times \pi_2X)$ splits as

$$\text{Wh}_1(\pi_1X; \mathbb{Z}_2 \times \pi_2X) = \text{Wh}_1(\pi_1X; \mathbb{Z}_2) \oplus \text{Wh}_1(\pi_1X; \pi_2X).$$

2.6.2 Construction of Θ

In this section we recall the definition of Θ set out in [HW73, Chapter VII]. Let f_t be a path in $\mathcal{F}(X)$ with endpoints in $\mathcal{E}(X)$. Suppose f_t lies in the kernel of Σ . We will associate to f_t an element of $\text{Wh}_1(\pi_1X; \mathbb{Z}_2 \times \pi_2X)$. In the quotient $\text{Wh}_1(\pi_1X; \mathbb{Z}_2 \times \pi_2X)/\chi(K_3\pi_1X)$ this element will be an invariant of the homotopy class in $(\mathcal{P}, \mathcal{E})$. This gives a map from $\ker \Sigma \subset \pi_1(\mathcal{F}, \mathcal{E})$, and hence from $\ker \Sigma \subset \pi_0(\mathcal{P})$, into $\text{Wh}_1(\pi_1(X); \mathbb{Z}_2 \times \pi_2(X))/\chi(K_3\pi_1X)$ which in both cases we denote by Θ .

By Section 2.4.4, after a deformation we may assume f_t is a collection of nested eyes, and that it has only handles of dimension i and $i + 1$ for $2 \leq i \leq n - 2$. As before we may assume that the births appear at times in $(0, \varepsilon)$ and the deaths occur at times in $(1 - \varepsilon, 1)$, and that index i critical points have critical value $< 1/2$ and that index $i + 1$ critical points have value $> 1/2$.

As previously we denote the middle level $f_t^{-1}(1/2)$ by V_t for $t \in [\varepsilon, 1 - \varepsilon]$, and label the belt spheres for the i -handles and attaching spheres for the $(i + 1)$ -handles by A_1, \dots, A_N and B_1, \dots, B_N in V_t respectively. In V_ε , $A_i \cap B_j$ consists of $\delta_{i,j}$ points. As we move forward in the time direction we see a homotopy of the A -spheres and B -spheres which keeps the B -spheres disjoint and embedded, keeps the A -spheres disjoint and embedded, but possibly introduces intersections between the A -spheres and the B -spheres. By general position we may assume in that in each t slice the intersection between A -spheres and B -spheres consists of disjoint double points. In $V_{1-\varepsilon}$ again we see A_i and B_j intersect in $\delta_{i,j}$ points. Note that there are no handle slides.

Consider the trace of this homotopy in the trace of the middle level. We adopt the notation $I_\varepsilon = [\varepsilon, 1 - \varepsilon]$, and denote

$$V \times I_\varepsilon = \bigcup_{t \in I_\varepsilon} V_t$$

for the trace of V_t . For the trace of the A -spheres we adopt the notation

$$A_i \times I_\varepsilon = \bigcup_{t \in I_\varepsilon} A_i$$

and for the trace of the B -spheres

$$B_i \times I_\varepsilon = \bigcup_{t \in I_\varepsilon} B_i.$$

The intersections of $A_i \times I_\varepsilon$ and $B_j \times I_\varepsilon$ for each i, j is a collection of lines and circles. In fact there is a single line of intersection in $A_i \times I_\varepsilon \cap B_i \times I_\varepsilon$, and all other intersection components are circles.

We will associate to each circle of intersection C , elements $\gamma_C \in \pi_1(X)$, $\sigma_C \in \pi_2(X)$ and $s_C \in \mathbb{Z}_2$. We construct a matrix M as follows

$$M_{i,j} = \sum_{\text{Circles of intersection } C \text{ between } A_i \text{ and } B_j} (\sigma_C + s_C)\gamma_C$$

Now $I + M$ is in $GL((\mathbb{Z}_2 \times \pi_2 X)[\pi_1 X])$; recall from Subsection 2.6.1 that

$$GL((\mathbb{Z}_2 \times \pi_2 X)[\pi_1 X]) = \{I + A \mid A \text{ has entries in } (\mathbb{Z}_2 \times \pi_2 X)[\pi_1 X]\}.$$

Hence $I + M$ specifies an element of $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$ as required.

Remark 2.6.6.

1. To motivate why these circles of intersection matter, if there were no circles and only a “straight line” of intersection in $A_i \times I_\varepsilon \cap B_j \times I_\varepsilon$, that is a line intersecting each t slice in a single point, then in fact we could remove this pair from the family and “cancel the eye”; that is deform the path f_t to one whose Cerf graphic did not contain this eye.
2. One might hope to remove a circle of intersection C between $A_i \times I_\varepsilon$ and $B_j \times I_\varepsilon$ by some ‘higher dimensional Whitney move’. That is, by finding discs $D_a \subset A_i \times I_\varepsilon$ and $D_b \subset B_j \times I_\varepsilon$ (the analogue of Whitney arcs), and finding an embedded 3-ball $B^3 \subset V \times I$ such that $\partial B^3 = D_a \cup D_b$ (the analogue of a Whitney disc). One would also ask that this 3-ball is framed, with framing on the boundary agreeing with a “Whitney Framing” specified by the discs D_a and D_b . One would also require that the intersection with any V_t contained no S^2 components. If we could find such a framed 3-ball we could perform an analogue of a Whitney move to remove the circle of intersection C . The element $\sigma_C \in \pi_2 X$ is an obstruction to finding such a ball, and the element $s_C \in \mathbb{Z}_2$ is the obstruction to correctly framing such a ball. In high dimensions the vanishing of these obstructions imply we can find and frame such a ball, and so perform a 3-dimensional Whitney trick to remove circles of intersection.

Construction of $\gamma_C \in \pi_1 X$

To obtain an element of $\pi_1 X$ associated to C , choose some time t such that $V \times t$ intersects C in some point p . pick an arc a from $*_{A_i}$ to p , and an arc b from $*_{B_j}$ to p . Then define

$$\gamma_C = \alpha_i \cdot a \cdot b^{-1} \cdot \beta_j^{-1}$$

in the same way that we defined the algebraic intersection in Section 2.1.3. To see independence from the choices of t , p , a and b , note that C is a circle in $A_i \times I_\varepsilon \cap B_j \times I_\varepsilon$,

and that both $A_i \times I_\varepsilon$ and $B_j \times I_\varepsilon$ are simply connected, and so C is null homotopic in both of these manifolds, and in $V \times I_\varepsilon$.

Construction of $\sigma_C \in \pi_2 X$

Since $A_i \times I_\varepsilon$ and $B_j \times I_\varepsilon$ are simply connected, the circle C bounds discs $D_a \subset A_i \times I_\varepsilon$ and $D_b \subset B_j \times I_\varepsilon$; note that these are possibly not embedded. We orient these discs by giving C an orientation; here we use the convention that at any point which is a positive intersection of A_i and B_j the orientation of C points in the positive t direction. Now the union $D_a \cup D_b$, along with the arc α_i (by convention) to the base point $*$ gives an element of $\pi_2 X$.

Note that both A_i and B_j are null-homotopic in X for all t (since they are belt spheres of 2-handles and attaching spheres of 3-handles), so $A_i \times I_\varepsilon$ and $B_j \times I_\varepsilon$ are null-homotopic in $X \times I_\varepsilon$ and so the above construction does not depend on the choice of D_a and D_b .

Construction of $s_C \in \mathbb{Z}_2$

First note that A_i and B_j intersect transversely, so we can identify the bundles

$$P := \nu(C, A_i \times I_\varepsilon) = \nu(B_j \times I_\varepsilon, V \times I_\varepsilon)|_C$$

This bundle is a $(n-2)$ -dimensional bundle over the circle. Since $A_j \times I_\varepsilon \cong S^{n-2} \times I$ is orientable, so is $\nu(C, A_j \times I_\varepsilon)$, so it must be the trivial bundle. In dimension 4 are \mathbb{Z} many choices of trivialisation (framing) for this bundle, in higher dimensions there are \mathbb{Z}_2 framings. The *difference* of any two choices of framing gives a well defined element of \mathbb{Z} in dimension 4 or \mathbb{Z}_2 in higher dimensions as in Remark 2.1.8. We will construct two framings of this bundle, and the difference (taking mod 2 when in dimension four) is $s_C \in \mathbb{Z}_2$.

Remark 2.6.7. One might ask, as Hatcher and Wagoner do in [HW73], if it is possible that in dimension 4 there is in fact a \mathbb{Z} valued invariant instead of a \mathbb{Z}_2 one. Igusa in [Igu21a] proves that there is such an integer valued invariant on the space of “marked lens space models”, that is paths in \mathcal{F} which have a single eye, with some additional “marking” information. He shows that when the marking information is dropped the \mathbb{Z} valued invariant only survives mod 2, suggesting a negative answer to this question; see [Igu21a, Lemma 1.12].

Two framings of P

For the first framing of P , consider $\nu(C, A_i \times I_\varepsilon)$; since C is null-homotopic in $A_i \times I_\varepsilon$, it has a canonical 0 -framing. We call this the A - 0 -framing.

For the second framing, we note that B_j is the attaching sphere for a handle, so $\nu(B_j, V)$ has a canonical framing that specifies how the handle is attached. Along with the isotopy of B_j this induces a framing of $\nu(B_j \times I_\varepsilon, V \times I_\varepsilon)$. Taking the restriction of this framing to $\nu(B_j \times I_\varepsilon, V \times I_\varepsilon)|_C$ gives the second framing. We call this the B -attaching-framing.

The difference of the A - 0 -framing and the B -attaching-framing up to homotopy gives an element of \mathbb{Z}_2 , (or in dimension 4 an element of \mathbb{Z} , which we then take mod 2 of) as in Remark 2.1.8.

2.6.3 Geometric description of 1-parameter families in $\ker \Sigma$ in dimension 4

We now describe some geometric features specific to dimension 4. We specify that X is a 4-manifold, with f_t as in Section 2.6.2.

Now we have that $V_t = X \#^N S^2 \times S^2$, and $A_1, \dots, A_n \in V_t$ and $B_1, \dots, B_n \in V_t$, are collections of 2-spheres. Again at $t = \varepsilon$ and $t = 1 - \varepsilon$. A_i intersects B_j in $\delta_{i,j}$ points, and at times in-between there is a homotopy of the spheres with possibly extra intersections between the A -spheres and the B -spheres. We can perturb this homotopy so that it can be seen as a sequence of finger moves and Whitney moves between the A -spheres and the B -spheres. It is a well known fact that any homotopy can be deformed so that all of the finger moves occur before the Whitney moves; this is a consequence of the fact that by dimensionality the guiding arcs for the finger moves may be freely isotoped whilst avoiding any Whitney discs, see [Qui86, Section 4.1] for a detailed treatment. We arrange that the finger moves occur before time $1/2$ and the Whitney moves occur after time $1/2$.

All of the data of this homotopy can now be seen in the level $V_{1/2}$; we call this copy of V the *middle-middle* level since it is $V_{1/2} = f_{1/2}^{-1}(1/2)$. At this time the spheres $A_i, B_j \in V_{1/2}$ intersect algebraically as $\delta_{i,j} \in \mathbb{Z}[\pi_1 X]$, but possibly have more geometric intersections. We see a family of framed embedded Whitney discs W_1, \dots, W_m which describe the Whitney moves that will be performed in the rest of the homotopy, and we also see framed embedded

Whitney discs U_1, \dots, U_m which undo the finger moves that were just performed. The Whitney discs U_1, \dots, U_m can be made disjoint from each other, as the finger moves are determined by arcs, so by dimensionality we can perturb these arcs to be disjoint from each other. Dually (considering the homotopy running backwards) we can arrange that the discs W_1, \dots, W_m are disjoint from each other. We refer to the finger move discs as the U -discs and the Whitney move discs as the W -discs. The U -discs may intersect the W discs.

Dual Spheres

There are also several dual spheres present. At time $t = \varepsilon$, $A_i \cap B_i = \{p\} \in V_\varepsilon$, and both A_i and B_i are embedded with trivial normal bundle; indeed they are $p \times S^2$ and $S^2 \times p$ fibers of an $S^2 \times S^2$ summand in $V_\varepsilon = X \#^N S^2 \times S^2$. Let A_i^* be a parallel copy of B_i , and let B_i^* be a parallel copy of A_i ; see Figure 2.9. Then A_i^* is a dual sphere to A_i since $A_i^* \cap A_i$ is a single point, $A_i^* \cap B_j = \emptyset$ for all j , and $A_i^* \cap A_j = \emptyset$ for $i \neq j$. Similarly for B_i^* . Note that A_i^* and B_j^* consists of $\delta_{i,j}$ points.




Figure 2.9: The spheres A_i and B_i , and the parallel copies which give transverse spheres A_i^* and B_i^* . We see that A_i^* intersects A_i in a single point, is disjoint from B_i , but intersects B_i^* in a single point; similarly B_i^* intersects B_i in a single point but is disjoint from A_i .

By dimensionality we can arrange that the arcs which determine the finger moves are disjoint from all dual spheres $A_1^*, \dots, A_n^*, B_1^*, \dots, B_n^*$. It follows that in the middle-middle level the spheres $A_1^*, \dots, A_n^*, B_1^*, \dots, B_n^*$ still intersect $A_1, \dots, A_n, B_1, \dots, B_n$ as above, and further are disjoint from the U -discs; they may however intersect the W -discs.

Noting that just before the deaths $A_i \cap B_j$ consists of $\delta_{i,j}$ points, we can repeat this argument running the homotopy backwards to obtain another set of dual spheres. Doing so we obtain (likely different) dual spheres $A_1^*, \dots, A_n^*, B_1^*, \dots, B_n^*$ in the middle-middle

level which are disjoint from the W -discs but not the U -discs. When we wish to emphasise the difference we will refer to the dual spheres disjoint from the U -discs as the *initial dual spheres*, and as the dual spheres disjoint from the W -discs as the *terminal dual spheres*.

2.6.4 Calculation for geometrically simple families

For X a 4-manifold we consider the following simple family of paths f_t in $\ker \Sigma$. First, a 2-handle 3-handle cancelling pair is created with spheres A and B in the middle level. Then, a single finger move is performed between A and B . Then, a single Whitney move is performed to remove the two intersections created by the finger move. Since A and B now intersect in a single point we then cancel them.

Considering the middle-middle level $V = V_{1/2}$, we may without loss of generality assume that the initial Whitney disc U (from the finger move) and the terminal Whitney disc W share the same Whitney arcs $\alpha \in A$ and $\beta \in B$. Indeed if they did not, we can isotope W to arrange that this is true as follows. Denote the Whitney arcs for W by $\alpha_W \subset A$ and $\beta_W \subset B$, and similarly $\alpha_U \subset A$, $\beta_U \subset B$ for the Whitney arcs of U . Note that A intersects B in three points, two of which are the endpoints of α_W . Since two arcs in a disc D^2 with the same endpoints are always isotopic, there is an isotopy taking α_W to α_U , which avoids the intersection points. We can extend this isotopy to the disc W in a small neighbourhood of A . We may do the same in a neighbourhood of B to make β_W agree with β_U .

The family f_t has a single circle of intersection C , so

$$\Theta(f_t) = ((s_C + \sigma_C)\gamma_C) \in \text{Wh}_1(\pi_1(X), \mathbb{Z}_2 \times \pi_2(X)).$$

It is also clear that $\gamma_C \in \pi_1(X)$ is the element of $\pi_1(X)$ associated to the finger move arc as in Remark 2.1.3.

To calculate σ_C we claim the following.

Proposition 2.6.8. *The element $\sigma_C \in \pi_2(X)$ can be represented by the 2-sphere $U \cup W$ joined to the basepoint $*$ by the basepoint arc for A , γ_A .*

To describe s_C we must work a little harder. Consider the bundle $\nu(B, V)|_\beta$. For either disc U and W we have 1-dimensional sub-bundles c_U, c_W of this bundle by taking the

intersection with the normal bundle for the disc,

$$c_U = \nu(B, V)|_\beta \cap \nu(U, V)|_\beta$$

$$c_W = \nu(B, V)|_\beta \cap \nu(W, V)|_\beta.$$

We arrange that the 1-bundles agree on $\nu(B, V)|_{\partial\beta} = \nu(\{p, q\}, A) = \{p, q\} \times D^2 \subset A$. Note that both U and W have Whitney arc α in A and so $c_U|_p$ and $c_W|_p$ are normal to α in $\nu(B, V)|_p = \nu(p, A) = p \times \mathring{D}^2$. Since in $\nu(p, A)$ there is only one choice of 1-bundle which is normal to $\alpha \subset \nu(p, A) = p \times \mathring{D}^2$ (up to isotopy), we can arrange that $c_U|_p = c_W|_p$. Similarly we can arrange that $c_U|_q = c_W|_q$. Arbitrarily we pick an orientation for these bundles to obtain two sections c_W and c_U of $\nu(B, V)|_\beta = S^2 \times \mathring{D}^2$ which agree on $\nu(B, V)|_{\partial\beta}$.

Remark 2.6.9. Note that these sections are the Whitney sections for W and U , but considered as a section of $\nu(B, V)|_\beta$ instead of $\nu(W, V)|_\beta$ and $\nu(U, V)|_\beta$. Recall that the definition of the Whitney section is asymmetric; we require the Whitney section for W be a section of $\nu(W, V)|_{\partial W}$ which is parallel to A on α and normal to B on β . Hence on β the Whitney section is precisely $c_W = \nu(B, V)|_\beta \cap \nu(W, V)|_\beta$.

Since c_W and c_U are sections of the same bundle and agree on the boundary we may consider the difference of these sections, which gives a well defined element of \mathbb{Z} as in Remark 2.1.8; note that Remark 2.1.8 refers to framings of $S^1 \times D^2$, here we consider framings of $I \times D^2$ which are fixed on $\partial I \times D^2$, which can also be seen to determine an element of $\pi_1 \text{GL}_2(\mathbb{Z}) = \mathbb{Z}$.

Proposition 2.6.10. *Consider the difference of c_U and c_W as an element of \mathbb{Z} as above. Taking this element of \mathbb{Z} mod 2 gives $s_C \in \mathbb{Z}_2$.*

We now proceed to prove both propositions.

Proof of Propositions 2.6.8 and 2.6.10. Consider $V \times I_\varepsilon$, say the finger move happens at time $\varepsilon \leq t_1 < 1/2$ and the Whitney move happens at time $1/2 < t_2 \leq 1 - \varepsilon$. We consider the sweep out of the Whitney disc W from time $1/2$ to time t_2 as the Whitney move is performed; considering the isotopy the Whitney disc takes as the Whitney move is performed, the sweep out is the union over all t of the Whitney disc. We denoted this sweep out by $Q_W \subset V \times [1/2, t_2] \subset V \times I_\varepsilon$. We similarly consider the sweep out of U , which we denote $Q_U \subset V \times [t_1, 1/2] \subset V \times I_\varepsilon$. See Figure 2.10 for a depiction of Q_U .




Figure 2.10: The ball Q_U swept out by the Whitney disc U through the time direction (denoted t).

Clearly Q_U is a ball, with boundary consisting of three discs, one of which is the disc $U \subset V \times \{1/2\}$. Another boundary component is the sweep out of the Whitney arc $\alpha_U \subset A$ denoted $D_\beta^U = Q_U \cap (A \times I_\varepsilon)$. The final boundary component is the sweep out of the Whitney arc β denoted $D_\alpha^U = Q_U \cap (B \times I_\varepsilon)$. We similarly denote the sweep out of $\alpha_W \subset A \times \{1/2\}$ by $D_\alpha^W \subset A \times [1/2, t_2]$ and sweep out of $\beta_W \subset A \times \{1/2\}$ by $D_\beta^W \subset A \times [1/2, t_2]$.

To prove Proposition 2.6.8, note that $D_\beta^W \cup D_\beta^U$ gives a disc in $B \times I_\varepsilon$ with boundary C , and $D_\alpha^W \cup D_\alpha^U$ gives a disc in $A \times I_\varepsilon$ with boundary C . Hence σ_C is represented by

$$(D_\alpha^W \cup D_\alpha^U) \cup (D_\beta^W \cup D_\beta^U)$$

along with the arc γ_A from A to the basepoint. Since $D_\alpha^U \cup D_\beta^U \cup U$ bounds Q_U , the disc $D_\alpha^U \cup D_\beta^U$ is isotopic to U . Similarly the disc $D_\alpha^W \cup D_\beta^W$ is isotopic to W via Q_W . Hence

$$(D_\alpha^W \cup D_\alpha^U) \cup (D_\beta^W \cup D_\beta^U) = (D_\alpha^U \cup D_\beta^U) \cup (D_\alpha^W \cup D_\beta^W) \simeq U \cup W$$

proving the Proposition 2.6.8.

To prove Proposition 2.6.10, define a 1-dimensional sub-bundle of $\nu(B \times I_\varepsilon, V \times I_\varepsilon)|_{D_\beta^U}$ via

$$S_U = \nu(B \times I_\varepsilon, V \times I_\varepsilon)|_{D_\beta^U} \cap \nu(Q_U, V \times I_\varepsilon)|_{D_\beta^U}$$

and a 1-dimensional sub-bundle of $\nu(B \times I_\varepsilon, V \times I_\varepsilon)|_{D_\beta^W}$ via

$$S_W = \nu(B \times I_\varepsilon, V \times I_\varepsilon)|_{D_\beta^W} \cap \nu(Q_W, V \times I)|_{D_\beta^W}.$$

Note that on β

$$S_U|_\beta = \nu(B \times I_\varepsilon, V \times I_\varepsilon)|_\beta \cap \nu(Q_U, V \times I)|_\beta = \nu(B, V)|_\beta \cap \nu(U, V)|_\beta = c_U$$

and similarly $S_W|_\beta = c_W$. As we did for c_W and c_U , we arrange that S_W and S_U agree on $\{p, q\} = \partial\beta$. We also pick orientations for S_W and S_U so that we can consider them as sections (picking the orientations so that they agree on $\{p, q\}$, and with c_U and c_W).

As $S_W|_C$ and $S_U|_C$ agree where they meet at p, q the union $S_W|_C \cup S_U|_C$ gives a section of $\nu(B \times I_\varepsilon, V \times I_\varepsilon)|_C = \nu(C, A \times I_\varepsilon) = S^1 \times D^2$; recall that this is the bundle which is used to define s_C . Consider the section S of $\nu(C, A \times I_\varepsilon)$ which points into the disc $D_\alpha^W \cup D_\alpha^U \subset A \times I_\varepsilon$. Clearly S is the 0-framed section of C in A (by definition). Since S_W is normal to Q_W , it is normal to $D_\alpha^W \subset Q_W$, and so $S_W|_C$ is normal to S . Similarly $S_U|_C$ is normal to S .

Hence $S_W|_C \cup S_U|_C$ is normal to S in $\nu(C, A \times I_\varepsilon)$. Since $\nu(C, A \times I_\varepsilon) = S^1 \times D^2$ this means that S and $S_W|_C \cup S_U|_C$ must be isotopic. Hence $S_W|_C \cup S_U|_C$ is the 0-framed section of C in $A \times I$. Hence the framing induced by $S_U|_C \cup S_W|_C$ is the A -0-framing of Section 2.6.2.

We divide the remainder of the proof into two cases.

Case 1: *on β the sections $c_W = S_W|_\beta$ and $c_U = S_U|_\beta$ agree up to isotopy.* In this case (after possibly an isotopy) $S_W \cup S_U$ is a section of $\nu(B \times I_\varepsilon, V \times I_\varepsilon)_{D_\beta^U \cup D_\beta^W}$. This section necessarily extends to a section of $\nu(B \times I_\varepsilon, V \times I_\varepsilon) = (S^2 \times I) \times D^2$. This section induces a trivialisation T of $\nu(B \times I_\varepsilon, V \times I_\varepsilon)$. Since $(S^2 \times I) \times D^2$ has a unique trivialisation, T must agree with the sweep out of the trivialisation used to attach the 3-handle.

Consider the trivialisation induced by $(S_W \cup S_U)|_C$ on $\nu(B \times I_\varepsilon, V \times I_\varepsilon)|_C$. This is necessarily the restriction of T to $\nu(B \times I_\varepsilon, V \times I_\varepsilon)|_C$, and hence is the restriction of the 3-handle framing to C . This is precisely the B -attaching-framing of Section 2.6.2, hence $(S_W \cup S_U)|_C$ induces the B -attaching-framing.

Since by the above $(S_W \cup S_U)|_C = S_W|_C \cup S_U|_C$ induces the A -0-framing, and the B -attaching framing we see that they agree and so $s_C = 0$ as required.

Case 2: *on β the sections $c_W = S_W|_\beta$ and $c_U = S_U|_\beta$ do not agree, but rather differ by n twists.* Given a trivial disc bundle $D^2 \times D^2 \rightarrow D^2$, a section of this bundle $s: D^2 \rightarrow D^2 \times D^2$, an arc $\gamma \subset \partial D^2$, and a section $r: \gamma \rightarrow \gamma \times D^2$, we can always perform an isotopy of the section s to obtain a new section s' such that $s|_\gamma = r$. Applying this to the disc bundle $\nu(B \times I_\varepsilon, V \times I_\varepsilon)|_{D_\beta^W}$ and the section S_W , we can perform an isotopy of S_W to obtain a




Figure 2.11: The operation of changing the section S_W so that it agrees with S_U on β . In the top left we depict $\nu(B \times I_\varepsilon, V \times I_\varepsilon)|_{D_\beta^W}$ and $\nu(B \times I_\varepsilon, V \times I_\varepsilon)|_{D_\beta^U}$, but for clarity only depict ∂D_β^W and ∂D_β^U , and $S_W|_{\partial D_\beta^W}$ and $S_U|_{\partial D_\beta^U}$. In the top left we see that the sections $S_U|_\beta$ and $S_W|_\beta$ differ by $n = 5$ twists. To obtain S'_W pictured on the right, we perform an isotopy which moves these twists to $\nu(B \times I_\varepsilon, V \times I_\varepsilon)|_{C \cap Q_W}$. Below we see that $S_W|_{C \cap Q_W} \cup S_U|_{C \cap Q_U}$ (bottom left) and $S'_W|_{C \cap Q_W} \cup S_U|_{C \cap Q_U}$ (bottom right) differ by $\pm n = \pm 5$ twists (depending on orientation).

section S'_W such that $S'_W|_\beta$ agrees with $S_U|_\beta$. We note that doing so changes $S'_W|_{C \cap Q_W}$. Since $S'_W|_\beta$ differs from $S_W|_\beta$ by n twists, it must be the case that $S'_W|_{C \cap Q_W}$ and $S_W|_{C \cap Q_W}$ differ by $\pm n$ twists (depending on the orientations), since the sections $S'_W|_{\partial W}$ and $S_W|_{\partial W}$ must agree up to isotopy. See Figure 2.11 for a depiction of this operation.

Repeating the above argument for when the framings agreed, using the framings S'_W and S_U , we see that the framing induced by $(S'_W \cup S_U)|_C$ is the B -attaching-framing. Again the framing induced by $S_W|_C \cup S_U|_C$ is the the A -0-framing. Since $(S'_W \cup S_U)|_C$ and $S_W \cup S_U$ differ by $\pm n$ twists, so do the B -attaching-framing and the A -0-framing, hence $s_C = \pm n \pmod 2 = n \pmod 2$ as required. \square

2.7 Realisation theorem for $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$

In dimension ≥ 5 ,

$$\Theta: \ker \Sigma \rightarrow \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \chi(K_3 \mathbb{Z}[\pi_1 X])$$

is surjective. Hatcher and Wagoner prove this when $k_1 X = 0$ [HW73, Chapter VII]; note that they claim to prove surjectivity for general $k_1 X$, however their definition of Θ is only valid when $k_1 X = 0$. The more general result was proved by Igusa [Igu84]. Hatcher and Wagoners proof considers $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$ as a quotient of

$$(\mathbb{Z}_2 \times \pi_2 X)[\pi_1 X] = \langle (s + \sigma)\gamma \mid \sigma \in \pi_2 X, s \in \mathbb{Z}_2, \gamma \in \pi_1 X \rangle.$$

They show that there exists a path $f_t \in \mathcal{F}$ with a single circle of intersection C such that $\gamma_C = \gamma$, $\sigma_C = \sigma$ and $s_C = s$ for arbitrary triples $\gamma \in \pi_1 X$, $\sigma \in \pi_2 X$, $s \in \mathbb{Z}_2$. In dimension 4 we cannot see a way to realise all possible triples.

Given a 4-manifold X , we denote the ‘‘Stiefel-Whitney classes’’ by $w_1^X: \pi_1(X) \rightarrow \{\pm 1\}$, and $w_2^X: \pi_2(X) \rightarrow \{0, 1\}$. These are the usual Stiefel-Whitney classes composed with the Hurewicz map and reduction mod 2, except that for w_1^X it will be useful for the target to be the multiplicative group of order two rather than the additive one, so if $w_1^X(\gamma)$ is trivial (usually 0) we send it to 1, if non-trivial (usually 1) we send it to -1 . The map w_1^X is often known as the ‘‘orientation character’’. We prove

Proposition 2.7.1. *Let X be a 4-manifold. Let $\sigma \in \pi_2 X$, $\gamma \in \pi_1 X$, and $s \in \mathbb{Z}_2$. If $w_2^X(\sigma) \neq 0$ or $s = 0$ then there exists a 1-parameter family f_t with a single circle of intersection C for which $\sigma_C = \sigma$, $\gamma_C = \gamma$, $s_C = s$.*

That is to say we can realise all elements $(\sigma + s)\gamma \in (\mathbb{Z}_2 \times \pi_2 X)[\pi_1]$ except for those with $w_2^X(\sigma) = 0$ and $s = 1$.

As a corollary we obtain Theorem C

Theorem C. *For X a compact 4-manifold, let*

$$\begin{aligned} \Xi &= \langle (s + \sigma)\gamma \mid w_2^X(\sigma) \neq 0 \text{ or } s = 0, s \in \mathbb{Z}_2, \sigma \in \pi_2 X, \gamma \in \pi_1 X \rangle \\ &\subset (\mathbb{Z}_2 \times \pi_2 X)[\pi_1 X] / \langle \alpha\gamma - \alpha^\tau \tau \gamma \tau^{-1}, \beta \cdot 1 \mid \alpha, \beta \in \mathbb{Z}_2 \times \pi_2 X, \tau, \gamma \in \pi_1 X \rangle \\ &= \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X). \end{aligned}$$

If $k_1 X = 0$ then $\Xi \subset \Theta(\ker \Sigma)$. Otherwise the same is true passing to the quotient $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \chi(K_3 \mathbb{Z}[\pi_1 X])$.

While we cannot realise all triples, we also cannot obstruct those exceptional values from

being realised. Indeed *topologically* one can find a homotopy of transverse spheres A and B with one circle of intersection realising all triples, as shown by [Kwa87].

We can however realise all values in $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$ by taking a single stabilisation with $S^2 \times S^2$.

Proposition 2.7.2. *Let X be a 4-manifold. Note that $\pi_1 X$ and $\pi_2 X$ include as subgroups of $\pi_1(X \# S^2 \times S^2)$ and $\pi_2(X \# S^2 \times S^2)$ respectively. For all values $\sigma \in \pi_2 X$, $\gamma \in \pi_1 X$ and $s \in \mathbb{Z}_2$ there exists a 1-parameter family $f_t: X \# S^2 \times S^2 \rightarrow I$ with a single circle of intersection C , such that $\sigma_C = \sigma$, $\gamma_C = \gamma$, $s_C = s$.*

Note that unlike the stable theorem for Σ , in this case, one $S^2 \times S^2$ stabilisation is enough. Indeed for any generator $(\sigma + s)\gamma$ we can come up with a 1-parameter family f_t and hence a corresponding pseudo-isotopy $F \in \mathcal{P}(X \# S^2 \times S^2)$ with $\Theta(F) = (\sigma + s)\gamma$. The composition of such pseudo-isotopies gives sums of generators without needing additional $S^2 \times S^2$ summands. This yields Theorem D.

Theorem D. *Let X be a compact 4-manifold. Note that $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$ includes in*

$$\text{Wh}_1(\pi_1(X \# S^2 \times S^2); \mathbb{Z}_2 \times \pi_2(X \# S^2 \times S^2)),$$

and identify $x \in \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$ with its image under this inclusion. There is a pseudo-isotopy F of $X \# S^2 \times S^2$, which is in $\ker \Sigma$ such that

$$\Theta(F) = x \in \text{Wh}_1(\pi_1(X \# S^2 \times S^2); \mathbb{Z}_2 \times \pi_2(X \# S^2 \times S^2)) / \chi(K_3 \mathbb{Z}[\pi_1(X \# S^2 \times S^2)]).$$

Note that unlike in the stable Σ case $\text{Wh}_1(\pi_1 X \# S^2 \times S^2; \mathbb{Z}_2 \times \pi_2(X \# S^2 \times S^2))$ is potentially a larger group, and there may be additional elements that we cannot realise (at least not without taking further connect sums).

We proceed with the proofs of both propositions.

Proof of Proposition 2.7.1. Fix $\sigma \in \pi_2 X$ and $\gamma \in \pi_1 X$. As in Remark 2.3.2, we build a path of Morse functions $f_t \in \mathcal{F}(X)$ by building a 1-parameter family of handle structures on $X \times I$, as in the proof of Theorem E.

We can represent σ by an immersed 2-sphere $S \subset X$ (with an arc η to the base point), and we can arrange that S intersects itself only at transverse double points.

We start with the trivial handle structure on $X \times I$. We then first create a cancelling 2-3 handle pair. When we do so, we ensure that the 2-3 handle creation takes place disjointly from $S \subset X$. We consider the middle level $V = X \# S^2 \times S^2$.

We denote the 2-handle belt sphere and 3-handle attaching sphere by A and B in V respectively; initially these are $p \times S^2$ and $S^2 \times q$ in $V = X \# S^2 \times S^2$. Note that we also have dual spheres A^* and B^* as in Section 2.6.3. Since we ensured the handle creation was disjoint from S , we see that S is disjoint from A , B , A^* , and B^* .

Since $A \cup B$ is π_1 -negligible in V , γ determines a finger move from A to B as in Remark 2.1.3. By dimensionality, we can ensure the finger move misses S and the dual spheres A^* and B^* . We perform this finger move to obtain a new handle structure. Afterwards, we see the Whitney disc U that undoes this finger move. The interior of U is disjoint from A , B , A^* and B^* .

Our plan now is to tube U into S to create a new Whitney disc W , then make W embedded and perform a Whitney move. We will do this so that $U \cup W$ represents σ and hence $\sigma_C = \sigma$, and so that W has the correct framing on β so that $s_C = s$. We deal with this in three cases.

Case 1: $w_2^X(\sigma) = 0$ and $s = 0$. First we perform interior twists to S so that $\nu(S, V)$ is the trivial disc bundle over the sphere. This is possible because $w_2^X(\sigma) = 0$. Now we tube U to S along an arc homotopic to $\gamma_A^{-1} \cdot \eta$ where γ_A is the arc from A to the base point; we ensure this arc misses any spheres. Since U and S are disjoint from A , B , A^* and B^* , so is W . Note that $W \cup U$ (with the basepoint arc γ_A) represents σ .

Since $\nu(S, V)$ is trivial, the framing of $\nu(U, V)$ extends to a framing of $\nu(W, V)$. Hence the disc-framing of U and the disc framing of W agree, so W is a framed (immersed) Whitney disc. We now remove each self intersection of W by pushing down into A as in Section 2.1.2. Pushing down each intersection one by one into A , W becomes embedded but now intersects A .

We may remove an intersection p between W and A by using the dual sphere A^* and the Norman trick; we take a path through A from p to $A \cap A^*$, then tubing W to a parallel copy of A^* using this path. Note that A^* is non-trivial in $\pi_2(V)$ but becomes trivial when included in $\pi_2(X \times I) = \pi_2(X)$, so this does not change the $\pi_2 X$ element given by $U \cup W$. We remove the intersections between W and A one by one, until W and A are disjoint.

Now W is a framed embedded Whitney disc, disjoint from A and B . We use W to perform a Whitney move. After this move the spheres intersect in a single point so we cancel the 2 and 3 handles. The resulting 1-parameter family f_t has a single circle of intersection C . By Section 2.6.4 we see that $\gamma_C = \gamma$. Since $U \cup W$ along with the base point arc γ_A represents σ , we have that $\sigma_C = \sigma$ by Section 2.6.4. We also see that the sections c_W and c_U of Section 2.6.4 agree; note that in a neighbourhood of the Whitney arc $\beta \subset B \subset V$, W and U agree by construction. Again by Section 2.6.4 we see that $s_C = 0$ as required.

Case 2: $w_2^X(\sigma) = 1$ and $s = 0$. In this case we cannot arrange $\nu(S, V)$ to be trivial, but we can perform interior twists to arrange it to be a bundle of Euler number 1 (note that an interior twist changes the Euler number of the normal bundle by two). We do this, then again tube U into S to form W . Doing so does not change the Whitney framing but does change the disc framing by a single twist, so W is not framed. We can fix this however by performing a single boundary twist between W and A . This makes W framed, but adds an intersection between A and W . We can now push down the self-intersections of W into A as before, then remove intersections between W and A (including the one from the boundary twist) using the dual sphere A^* .

As in Case 1 W is now embedded so we perform a Whitney move, then cancel the handles. Again, we see that in a neighbourhood of $\beta \subset B \subset V$, W and U agree (since we boundary twisted with A), so again the sections c_W and c_U agree, so as in Case 1 we have that $s_C = 0$. It similarly follows that $\gamma_C = \gamma$ and $\sigma_C = \sigma$.

Case 3: $w_2^X(\sigma) = 1$ and $s = 1$. As in Case 2 we interior twist so that $\nu(S, V)$ has normal bundle one, then tube U to S to form W . To make W framed we now perform a single boundary twist with B .

We now push down the self intersections of W into B , then resolve the intersections of W and B using B^* to make W embedded. Again we now perform the Whitney move and cancel the handles.

In the neighbourhood of β , U and W differ by a single twist (from the boundary twist we performed). Hence c_W and c_U also differ by a single twist. Hence $\gamma_C = \gamma$, $\sigma_C = \sigma$ and $s_C = 1$ by Section 2.6.4. \square

Before we prove the stable version of the theorem, it is useful to see what goes wrong when we try to realise $s_C = 1$ when $w_2^X(\sigma) = 0$. Using the notation from the proof of Proposition

2.7.1 we consider the Whitney disc U which undoes the finger move and tube it to S , which we arranged to have trivial normal bundle, to form a new Whitney disc W . Since we would like that c_W and c_U do not agree (mod 2), we would like to do an odd number of boundary twists of W with B . However doing so unframes the disc, and this cannot be fixed using interior twists (since these change the disc framing by an even number). Hence we must perform at least one boundary twist of W with A . However if we do then W intersects both A and B . This is a problem as the dual spheres A^* and B^* intersect; we cannot use the Norman trick on both A^* and B^* without creating new self intersections of W .

We can however make the above construction work in the case that we have an $S^2 \times S^2$ summand. We will use this summand to build an additional dual sphere, which can be used to resolve the intersections.

Proof of Proposition 2.7.2. The only case left to realise is when $w_2^X(\sigma) = 0$ and $s = 1$. Consider $X \# S^2 \times S^2$, let $Y = S^2 \times S^2$, let $X' = X \# Y$ and let $\widehat{X} = X \setminus B^4 \subset X'$.

Given $\sigma \in \pi_2 X$, we again represent σ by an immersed sphere $S \subset \widehat{X}$ with intersection points transverse double points; note that $\pi_2 \widehat{X} = \pi_2 X$. We again arrange that $\nu(S, \widehat{X})$ is the trivial disc bundle over S^2 by performing interior twists to S .

We again construct a 1-parameter family f_t for the 4-manifold $X \# S^2 \times S^2$. Following the steps in the proof of Proposition 2.7.1 we first create a cancelling 2-3 handle pair in $\widehat{X} \times I$, and consider the middle level of $\widehat{X} \times I$, denoted $\widehat{V} = \widehat{X} \# S^2 \times S^2$ which is a subset of the middle level of $X \times I$, denoted $V = X \# Y \# S^2 \times S^2$. Again we ensure that this 2-3 handle pair is created in the complement of S . To distinguish this copy of $S^2 \times S^2$, we let $Z = S^2 \times S^2$ and say $\widehat{V} = \widehat{X} \# Z$, and note that $V = \widehat{V} \cup (Y \setminus B^4)$

We again denote the 2-handle belt sphere and 3-handle attaching spheres by A and B respectively. We now perform a finger move, in \widehat{V} , corresponding to $\gamma \in \pi_1 X = \pi_1 \widehat{X}$, ensuring this finger move misses A , B , A^* , and B^* . In \widehat{V} we again see a Whitney disc U that undoes the move. We again tube U to S to create an immersed framed Whitney disc W .

To obtain the correct value for s_C , we now perform a single boundary twist of W with B . We then perform another boundary twist on W with A (an opposite twist), so that W remains a framed Whitney disc. However W now intersects A and B . We resolve the

intersection single between W and A using the Norman trick on A^* . We then resolve the single intersection with B using the Norman trick on B^* , noting that this adds a single further self-intersection of W as the parallel copy of A^* intersects the parallel copy of B^* . We consider the Clifford torus $T \subset \widehat{V}$ for the Whitney disc W . The torus T intersects W in exactly one point, but is disjoint from A, B, A^* and B^* . See Figure 2.12 for a depiction of this torus.




Figure 2.12: We present 4-dimensional space as slices of 3-dimension space. In the middle slice we see the Whitney disc W and a disc subset of the sphere B . In each slice we see a line from A , which sweep out a disc subset from A . We picture the Clifford torus T , and see that it intersects W in exactly one point, but does not intersect A or B . In the fourth frame we depict the cap D_A , which intersects A in exactly one point. To see D_B one can draw a new picture with the roles of A and B reversed.

There are two caps for T , namely embedded discs D_A and D_B , which intersect T exactly on ∂D_A and ∂D_B ; we picture D_A in the fourth frame in Figure 2.12. The disc D_A intersects A in exactly one point but is disjoint from W, B, A^* and B^* , and the disc D_B intersects B in exactly one point but is disjoint from W, A, A^* and B^* . Further D_A and D_B intersect only in a single point on their boundary; see Figure 2.13.




Figure 2.13: The Clifford torus with the caps. Note that A and B do not intersect T , and that W only intersects T in a single point.

We now manipulate the caps. We direct the reader to Figure 2.16 for this manipulation. We remove the intersection of D_A and A using the Norman trick, tubing it into A^* . We similarly remove the intersection of D_B and B using B^* , noting that this adds intersections with W , and with the dual spheres A^* and B^* , and adds a single point of intersection




Figure 2.14: The symmetric capping operation. In the top picture we see T and the caps D_A and D_B . In the bottom picture we see the two parts of the symmetric capping; left we see T with the neighbourhoods of ∂D_A and ∂D_B removed, and right the parallel copies of the caps, with the extra square around $\partial D_A \cap \partial D_B$ which we glue back in; we have highlighted this square glued back in in red.

between D_A and D_B . See the first picture of Figure 2.16 for a depiction of the resulting discs.

We can now perform a “symmetric capping” operation of Freedman and Quinn [FQ90] using the two discs D_A and D_B . To do this, we remove the neighbourhood of $\partial D_A \cup \partial D_B$ from T (Noting these circles intersect in a single point) and glue back in two parallel copies of D_A , two parallel copies of D_B , and we glue back in a square in T around $\partial D_A \cap \partial D_B \in T$ to fill the resulting hole; see Figure 2.14. We smooth the edges of the resulting sphere and denote the result of this operation by P .

Note that since D_A and D_B have a single point of intersection, P is an immersed sphere with four double points. P is disjoint from D_A and D_B , but since the caps intersected W , P has additional points of intersection with W . See the top left and bottom left pictures of Figure 2.15; in the top left we see an intersection of D_B and W , in the bottom left we see how this introduces two points of intersection between W and P .

We may however remove these new intersections of W and P in the following way. First, consider the intersections of D_B and W . We may push down all such intersections into T so that the intersection points of W and T lie on ∂D_A ; see the top right picture in Figure 2.15. Now when we take the symmetric capping to obtain P , W will not intersect P ;

see the bottom right picture in Figure 2.15. We now do the same thing for intersections between W and D_A . Note that this operation introduces extra self-intersections of W ; indeed each pair of intersections $p \in D_A \cap W, q \in D_B \cap W$ gives rise to a single self-finger move of W , which is performed when we push these intersections down into T . See the second and third picture of Figure 2.16 for another depiction of the symmetric capping operation.




Figure 2.15: Removing intersections with S which arise from the intersections with the caps. In the top figure we see T , the caps D_A and D_B and a point of intersection between U and D_B (note that W continues into the past and future). In the bottom left we see the two parts of the symmetric capping P ; we draw them separately to make the picture clear. We picture W in both these parts (we draw W twice to show how it interacts with both parts of the symmetric capping). In the top right figure we see the result of pushing down the intersection of W and D_B into T . In the bottom left we see that W now misses the symmetric capping.

After this process, W and P intersect only in a single point (the point in which the original Clifford torus intersected W). W has many self intersections, but still does not intersect A or B in its interior. Further, P is disjoint from A and B . P is immersed with 4 double points which arose from the point of self intersection between D_A and D_B . Hence we can see the intersection points of P as the intersection of two parallel copies of D_A , denoted D_A^+ and D_A^- , and two parallel copies of D_B , denoted D_B^+ and D_B^- ; see the third picture of Figure 2.16.

Since so far we worked entirely within \widehat{V} , we may now use the $S^2 \times S^2$ summand

$$Y \setminus B^4 = S^2 \times S^2 \setminus B^4 \subset V = \widehat{V} \cup (Y \setminus B^4)$$

to resolve the self-intersections of P . First, we take the parallel sheets D_A^+ and D_A^- and




Figure 2.16: We depict the manipulation of the Clifford caps. To obtain the first picture we tube D_A into A^* and D_B into B^* . To obtain the second picture from the first we push down the intersections between W and D_A into T (we do the same for any intersections between W and D_B). To obtain the third picture we perform the symmetric capping operation, and note that W is now disjoint from the discs, and from $T \setminus (\partial D_A \cup \partial D_B)$. In the third picture we see the sphere P as the union of $T \setminus (\partial D_A \cup \partial D_B)$, the caps, and the square around $\partial D_A \cap \partial D_B$. To obtain the fourth picture we perform a two sheeted Norman trick, tubing D_A^\pm into $p \times S^2 \subset Y \setminus B^4$. To obtain the final picture we perform a further two sheeted Norman trick, tubing D_B^\pm into $S^2 \times p \subset Y \setminus B^4$. In the final picture we see P' as the union of $T \setminus (\partial D_A \cup \partial D_B)$, the caps, and the square around $\partial D_A \cap \partial D_B$. We note that P' is embedded with a single point of intersection with W .

tube them to two parallel copies of $p \times S^2 \subset Y \setminus B^4$; see Figure 2.17. We tube using some arc which we make disjoint from any spheres and discs by dimensionality. See the fourth picture of Figure 2.16 for a depiction of the resulting discs. Another way to see this operation is to consider tubing D_A into $p \times S^2$, then taking two parallel copies of the resulting surface; hence D_A^\pm are still parallel copies of some surface.




Figure 2.17: Tubing the discs D_A^+ and D_A^- into two parallel copies of $p \times S^2$.

The two parallel sheets D_A^\pm now intersect $S^2 \times q$. We can remove the intersections of D_A^\pm with D_B^\pm by performing a two-sheeted Norman trick, tubing the two sheets into $S^2 \times q$ to remove all self-intersections of P . See Figure 2.18 for a depiction of this. See the fifth Figure of 2.16 for a depiction of the resulting discs. Another way to see this is to view D_A^\pm as parallel copies of D_A which intersects $S^2 \times q$ in a single point, and D_B^\pm as parallel copies of D_B . Performing the Norman trick with $S^2 \times q$ to remove the single point of intersection between D_A and D_B , then taking parallel copies of both yields the same result. We denote the resulting surface by P' . The sphere P' is embedded, intersects W in a single point, and is disjoint from A and B .

Note that we could have equivalently used $Y \setminus B^4 = S^2 \times S^2 \setminus B^4$ to make D_A and D_B disjoint, then taken parallel copies of both.

Since D_A^+ and D_A^- had opposite orientations, and so did D_B^+ and D_B^- , P' represents the same homotopy class in V as P (this is precisely why we went to so much trouble with the symmetric capping). Since P is contained in the $S^2 \times S^2$ summand of $\widehat{V} = \widehat{X} \# S^2 \times S^2$, when we include P into $X \times I$, it is 0 in $\pi_2(X \times I) = \pi_2 X$, hence $[P'] = 0 \in \pi_2(X \times I)$.

Since P' is embedded with trivial normal bundle (note that all the spheres we tubed into had trivial normal bundle), and P' intersects W in a single point, we can use P' to perform the Norman trick and remove the self intersections of W . After doing so W is a framed




Figure 2.18: Left we see D_A^+ and D_A^- after we tube them into parallel copies of $p \times S^2$. We also see the four intersection points with D_B^+ and D_B^- . Right we perform a two sheeted Norman trick, tubing D_B^+ and D_B^- into parallel copies of $S^2 \times q$, removing the intersections.

embedded Whitney disc, and does not intersect A or B in its interior. Note also, since P' is 0 in $\pi_2 X$, the union of the Whitney discs $U \cup W$ (with the basepoint arc) still represents σ in $\pi_2 X$.

We now use W to perform a Whitney move, then cancel the resulting handles. As previously we see that $\gamma_C = \gamma$, $\sigma_C = \sigma$ and $s_C = 1$, the latter because we performed a single boundary twist with B . □

2.8 Pseudo-isotopy versus isotopy

In this section we construct diffeomorphisms of 4-manifolds which are pseudo-isotopic but not isotopic to the identity.

To do this we make use of the previously described obstructions Θ and Σ . Given a diffeomorphism $f \in \text{Diff}_{PI}(X, \partial X)$, we may try to obstruct it from being isotopic to the identity by picking a pseudo-isotopy $F \in \mathcal{P}$ from the identity to f , that is with $F|_{X \times 1} = f$, and then evaluating $\Sigma(F)$ and $\Theta(F)$. However, $\Sigma(F)$ and $\Theta(F)$ clearly depend on the choice of F , not just on f . Given another choice of pseudo-isotopy $G \in \mathcal{P}$ from the identity to f , we see that the composition $F \circ G^{-1}$ is a pseudo-isotopy fixing the entire boundary of $X \times I$. This motivates the following definition.

Definition 2.8.1. Let $F \in \mathcal{P}(X)$ be a pseudo-isotopy with $F_{X \times 1} = \mathbf{1}_X$. Note that $F_{X \times 0}$ and $F_{\partial X}$ are the identity by definition of \mathcal{P} so in fact $F_{\partial(X \times I)} = \mathbf{1}_{\partial(X \times I)}$. We say that F is

an *inertial* pseudo-isotopy, and denote the set of inertial pseudo-isotopies by $\mathcal{J}(X) \subset \mathcal{P}(X)$, or just \mathcal{J} when X is clear from the context.

Remark 2.8.2. Since inertial pseudo-isotopies fix the entire boundary of $X \times I$ in fact

$$\mathcal{J}(X) = \text{Diff}(X \times I, \partial(X \times I)).$$

Let $f: X \rightarrow X$ be a diffeomorphism that is pseudo-isotopic to the identity. Define

$$\Sigma: \pi_0 \text{Diff}_{PI}(X, \partial X) \rightarrow \text{Wh}_2(\pi_1 X) / \Sigma(\mathcal{J})$$

by saying $\Sigma(f) := \Sigma(F)$, where F is some pseudo-isotopy from $\mathbb{1}_X$ to f ; by the discussion above this is independent of the choice of F .

To show $\Sigma(f)$ is an invariant of the isotopy class of f consider g isotopic of f , and let S be an isotopy from the identity to $f^{-1} \circ g$. Pick a pseudo-isotopy F from the identity to f , then $F \circ S$ is a pseudo-isotopy from the identity to g , so

$$\Sigma(f) = \Sigma(F) = \Sigma(F) + \Sigma(S) = \Sigma(F \circ S) = \Sigma(g).$$

Note here that $\Sigma(S) = 0$ since S is an isotopy.

We wish to also define $\Theta(f) \in \left(\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \chi(K_3 \mathbb{Z}[\pi_1 X]) \right) / \Theta(\mathcal{J} \cap \ker \Sigma)$. To define Θ for $f \in \ker \Sigma$ we first prove the below lemma.

Lemma 2.8.3. *Let f be a self-diffeomorphism of X which is pseudo-isotopic to the identity. If $\Sigma(f) = 0 \in \text{Wh}_2(\pi_1 X) / \Sigma(\mathcal{J})$, then there exists a pseudo-isotopy F from $\mathbb{1}_X$ to f such that $\Sigma(F) = 0 \in \text{Wh}_2(\pi_1 X)$.*

Proof. Let G be a pseudo-isotopy from the identity to f . Since $\Sigma(f) = 0 \in \text{Wh}_2(\pi_1 X) / \Sigma(\mathcal{J})$, we have $\Sigma(G) \in \Sigma(\mathcal{J})$. Hence there exists $S \in \mathcal{J}$ with $\Sigma(G) = \Sigma(S)$. Let $F = G \circ S^{-1}$, since S is inertial F is also a pseudo-isotopy from f to the identity. Then $\Sigma(F) = \Sigma(G \circ S^{-1}) = \Sigma(G) - \Sigma(S) = 0$ as required. \square

Hence given $f \in \ker \Sigma$, Taking F to be a pseudo-isotopy from the identity to f with $\Sigma(F) = 0$, we define.

$$\Theta: \ker \Sigma \subset \pi_0 \text{Diff}_{PI}(X, \partial X) \longrightarrow \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \chi(K_3 \mathbb{Z}[\pi_1 X]) / \Theta(\mathcal{J} \cap \ker \Sigma)$$

$$f \longmapsto \Theta(F).$$

Note that the $\ker \Sigma$ on the right refers to the subset of \mathcal{P} , while the $\ker \Sigma$ on the left refers to the subset of $\text{Diff}_{PI}(X, \partial X)$.

To see that this is well defined, suppose f and g are isotopic, and let $F, G \in \ker \Sigma$ be pseudo-isotopies from the identity to f and g respectively as in Lemma 2.8.3. Also let S be an isotopy from the identity to $f^{-1} \circ g$. Then

$$\Theta(f) - \Theta(g) = \Theta(F) - \Theta(G) = \Theta(F) + \Theta(S) - \Theta(G) = \Theta(F \circ S \circ G^{-1}) \in \Theta(\mathcal{J} \cap \ker \Sigma)$$

If $k_1 X = 0$ we may also define $\Theta(f)$ when $\Sigma(f) \neq 0$, however in order to obtain something well defined we must define it in the group $\Theta(f) \in \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \Theta(\mathcal{J})$. Note that we possibly quotient out by a larger subgroup. We conjecture that these groups are the same.

Conjecture 2.8.4. *For X a 4-manifold with $k_1 X = 0$, $\Theta(\mathcal{J}(X)) = \Theta(\mathcal{J}(X) \cap \ker \Sigma)$*

The images $\Theta(\mathcal{J})$ and $\Sigma(\mathcal{J} \cap \ker \Sigma)$ are in general difficult to determine. For 4-manifolds of a certain form however we can say something about these two subgroups, in particular when $X = M^3 \times [0, 1]$ where M is a 3-manifold.

2.8.1 Duality formulae

We recall the duality formulae and involutions of the Whitehead groups from Chapter VIII of [HW73] and Sections 4 and 5 of [Hat73]. We first give an analogue of turning a Morse function upside down for pseudo-isotopies.

Definition 2.8.5. Let F be a pseudo-isotopy. Denote the reflection map on $X \times I$ which sends (p, s) to $(p, 1 - s)$ by R . We define the *dual* pseudo-isotopy to F to be

$$\bar{F} = \left((F|_{X \times 1})^{-1} \times \mathbf{1}_I \right) \circ R \circ F \circ R.$$

We note \bar{F} sends $X \times i$ to $X \times i$ for $i \in \{0, 1\}$ and that $(R \circ F \circ R)|_{X \times 0} = F|_{X \times 1}$, so

$$\bar{F}|_{X \times 0} = \left(\left((F|_{X \times 1})^{-1} \times \mathbf{1}_I \right) \circ R \circ F \circ R \right) |_{X \times 0} = F|_{X \times 1}^{-1} \circ F|_{X \times 1} = \mathbf{1}_X.$$

Further $(R \circ F \circ R)|_{\partial X \times I} = \mathbf{1}_{\partial X \times I}$ so $\bar{F}|_{\partial X \times I} = \mathbf{1}_{\partial X \times I}$, so $\bar{F} \in \mathcal{P}(X)$.

Let f_t be a path in \mathcal{F} representing F ; that is with $f_1 = p \circ F$, $f_0 = p$. Denote

$$\bar{f}_t = R_I \circ f_t \circ R = 1 - f_t \circ R,$$

where R_I is the reflection map on I . It is clear that \bar{f}_t is a 1-parameter family for \bar{F} . Considering the index of the Morse critical points, we see that a critical point of index i becomes a critical point of index $n - i$.

One can view \bar{f}_t as turning each Morse function upside down (which is usually considered to be taking $R_I \circ f$ for a Morse function f), then using the diffeomorphism R to fix the ends, that is to make $\bar{f}_t(X \times i) = i$ for $i \in \{0, 1\}$. Because R is a diffeomorphism it does not change the index of the critical points.

Hatcher in [Hat73] in fact uses the path $1 - f_t$ to compute $\Sigma(\bar{F})$ and $\Theta(\bar{F})$. Again because $R_{X \times I}$ is a diffeomorphism it does not change the index of handles, or the intersections between handles, and it is the identity on homotopy groups, so does not change the computation. We describe this computation below.

Involution of $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$

In this section we recall the involution

$$\begin{aligned} \bar{\cdot}: \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) &\longrightarrow \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) \\ \theta &\longmapsto \bar{\theta} \end{aligned}$$

described by Hatcher in [Hat73, Lemma 4.3]. We will use this to relate $\Theta(\bar{F})$ and $\Theta(F)$.

We define maps

$$\begin{aligned} \bar{\cdot}: \mathbb{Z}[\pi_1 X] &\longrightarrow \mathbb{Z}[\pi_1 X] \\ \gamma &\longmapsto w_1^X(\gamma)\gamma^{-1} \text{ for } \gamma \in \pi_1 X \end{aligned}$$

and

$$\begin{aligned} \bar{\cdot}: (\mathbb{Z}_2 \times \pi_2 X)[\pi_1 X] &\longrightarrow (\mathbb{Z}_2 \times \pi_2 X)[\pi_1 X] \\ (n, \sigma)\gamma &\longmapsto \left(n + w_2^X(\sigma), -w_1^X(\gamma)\sigma\gamma^{-1} \right) \gamma^{-1}. \end{aligned}$$

For $\gamma \in \pi_1 X$, $\sigma \in \pi_2 X$, $n \in \mathbb{Z}_2$. Here $\sigma^{\gamma^{-1}}$ denotes the action of γ^{-1} on σ . Clearly this defines these maps on the whole group by additivity.

We can define an involution on $GL((\mathbb{Z}_2 \times \pi_2 X)[\pi_1 X])$

$$I + A_{i,j} \mapsto I + \overline{A_{j,i}}.$$

This in turn induces an involution on $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$, which we denote $\theta \mapsto \bar{\theta}$. Considering the isomorphism in Corollary 2.6.4, we can equivalently define the involution on $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$ by considering it as a quotient of $(\mathbb{Z}_2 \times \pi_2 X)[\pi_1 X]$, then the involution is induced directly by the involution on $(\mathbb{Z}_2 \times \pi_2 X)[\pi_1 X]$.

Remark 2.8.6. Note that this involution depends not only on the groups $\pi_1 X$ and $\pi_2 X$ and the action of $\pi_1 X$ on $\pi_2 X$ (which is all that is needed to define the Whitehead groups) but also on the Stiefel-Whitney classes w_1^X and w_2^X .

We also wish to define an involution on the target of Θ when $k_1 X \neq 0$, namely the quotient $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)/\chi(K_3 \mathbb{Z}[\pi_1 X])$, however in order to do so we would need the involution to fix $\chi(K_3 \mathbb{Z}[\pi_1 X])$; we conjecture that this is true.

Conjecture 2.8.7. $\overline{\chi(K_3 \mathbb{Z}[\pi_1 X])} = \chi(K_3 \mathbb{Z}[\pi_1 X])$.

We suspect that one can define an involution on $K_3 \mathbb{Z}[\pi_1 X]$ so the involution commutes with χ , which would prove this conjecture, however we have not been able to resolve this. In order to avoid this problem, instead we define

$$\hat{\chi} := \chi(K_3 \mathbb{Z}[\pi_1 X]) + \overline{\chi(K_3 \mathbb{Z}[\pi_1 X])}$$

and

$$\hat{\Theta}: \pi_0 \mathcal{P}(X) \longrightarrow \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)/\hat{\chi}$$

by taking the composition of Θ and the quotient map

$$\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)/\chi \rightarrow \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)/\hat{\chi}.$$

Now we can also define

$$\widehat{\Theta}: \pi_0 \text{Diff}_{PI}(X, \partial X) \longrightarrow (\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \widehat{\chi}) / \widehat{\Theta}(\mathcal{J} \cap \ker \Sigma)$$

as we did for Θ .

Clearly when $k_1 X = 0$ or $\chi(K_3 \mathbb{Z}[\pi_1 X]) = 0$, $\Theta = \widehat{\Theta}$, and if Conjecture 2.8.7 holds then $\Theta = \widehat{\Theta}$ regardless.

Involution of $\text{Wh}_2(\pi_1 X)$

We now recall the involution of $\text{Wh}_2(\pi_1 X)$ defined by Hatcher-Wagoner; see [HW73, Chapter VIII]. We use this to relate $\Sigma(\overline{F})$ and $\Sigma(F)$.

First we define an involution on the Steinberg group $\text{St}(\mathbb{Z}[\pi_1 X])$ group denoted $a \mapsto \bar{a}$. We define the map on the generators by $s_{i,j}^\lambda \mapsto s_{j,i}^{\bar{\lambda}}$; note that $\lambda \in \mathbb{Z}[\pi_1]$ and that we use the involution defined in Section 2.8.1. Since this preserves the Steinberg relations it defines an involution of $\text{St}(\mathbb{Z}[\pi_1 X])$.

One can easily define an involution on $E(\mathbb{Z}[\pi_1 X])$ sending $M_{i,j}$ to $\overline{M_{j,i}}$, and it is clear that this commutes with the map $\pi: \text{St}(\mathbb{Z}[\pi_1 X]) \rightarrow E(\mathbb{Z}[\pi_1 X])$, which means the involution is defined on $K_2(\mathbb{Z}[\pi_1 X])$.

Since $\overline{w_{i,j}^{\pm g}} = w_{j,i}^{\pm \bar{g}}$ the involution sends $W(\pm \pi_1 X)$ to itself; recall that $W(\pm \pi_1 X)$ is the subgroup of $\text{St}(\mathbb{Z}[\pi_1 X])$ defined in Section 2.4.1, and that $\text{Wh}_2(\pi_1[X]) = K_2(\mathbb{Z}[\pi_1[X]]) / W(\pm \pi_1 X)$.

Hence it follows that the involution is defined in the quotient $\text{Wh}_2(\pi_1[X])$ as required.

Duality formulae

We can now state the duality formulae of Hatcher and Wagoner.

Proposition 2.8.8 ([HW73, Chapter VIII], [Hat73, Duality Formula 4.4]). *Let X be a manifold of dimension n . Then*

$$\Sigma(\overline{F}) = (-1)^n \overline{\Sigma(F)}$$

and if $F \in \ker \Sigma$

$$\widehat{\Theta}(\overline{F}) = (-1)^n \overline{\widehat{\Theta}(F)},$$

where on the right hand side we use the involutions defined in Sections 2.8.1 and 2.8.1.

Hatcher and Wagoner prove this for Θ which is defined when $k_1X = 0$, however taking the additional quotient by $\widehat{\chi}$ as in Section 2.8.1 makes no difference to the proof. Indeed for Θ they consider the change of path of Morse functions from f_t to $1 - f_t$, and compare the change on the elements σ_C, γ_C, s_C for each circle of intersection. They show that it corresponds to the involution we defined on $(\mathbb{Z}_2 \times \pi_2X)[\pi_1X]$, then pass to the quotient $\text{Wh}_1(\pi_1X; \mathbb{Z}_2 \times \pi_2X)$; passing to a further quotient by $\widehat{\chi}$ does not change the argument.

2.8.2 Inertial pseudo-isotopies of $M^3 \times I$

When $X = M \times I$ is the product of a 3-manifold M and the interval, we can say more about the image of \mathcal{J} under $\widehat{\Theta}$ and Σ .

We first note that there is a differential defined on $\text{Wh}_2(\pi_1X) \oplus \text{Wh}_1(\pi_1X; \mathbb{Z}_2 \times \pi_1X)/\widehat{\chi}$ given by $d_i(x) = x - (-1)^i \bar{x}$. Note that this is defined independently on the summands. We define $Z_i = \ker d_i$ and $B_i = \text{im } d_{i+1}$, we will also split these out as

$$B_i = B_i^2 \oplus B_i^1 \subset \text{Wh}_2(\pi_1X) \oplus \text{Wh}_1(\pi_1X; \mathbb{Z}_2 \times \pi_1X)/\widehat{\chi}$$

and

$$Z_i = Z_i^2 \oplus Z_i^1 \subset \text{Wh}_2(\pi_1X) \oplus \text{Wh}_1(\pi_1X; \mathbb{Z}_2 \times \pi_1X)/\widehat{\chi}.$$

We recall the result of Hatcher that will allow us to bound the size of $\Sigma(\mathcal{J})$ and $\Theta(\mathcal{J} \cap \ker \Sigma)$ for $M \times I$.

Proposition 2.8.9. [Hat73, Lemma 5.3] *Let M be an $(n - 1)$ -manifold, let $X = M \times I$, and let $F \in \mathcal{J}(X)$ be an inertial pseudo-isotopy. Then*

$$\Sigma(\mathcal{J}) = (-1)^n \overline{\Sigma(\mathcal{J})}$$

and if $J \in \ker \Sigma$,

$$\widehat{\Theta}(J) = (-1)^n \overline{\widehat{\Theta}(J)}.$$

Hence we have

$$\Sigma(\mathcal{J}) \subset Z_n^2$$

and

$$\Theta(\mathcal{J} \cap \ker \Sigma) \subset Z_n^1.$$

Hatcher proves this for Θ but it also holds for $\widehat{\Theta}$ when $k_1 X \neq 0$. We recall the proof below.

Proof of Proposition 2.8.9. Let R denote the map on $X \times I$ sending (x, s) to $(x, 1 - s)$. Since $X = M \times I$, there is also an involution on X sending (m, l) to $(m, 1 - l)$, which in turn induces an involution on $X \times I = M \times I \times I$, which we denote by L .

We can define a further map on $M \times I \times I$ by rotating around the I^2 factor; we denote this rotation R_θ for $\theta \in [0, 2\pi]$. Define

$$\widetilde{J} = R_\pi \circ J \circ R_\pi.$$

Noting that $R_\pi = R \circ L$ we have that $\widetilde{J} = L \circ \overline{J} \circ L$. Since conjugation by L induces the identity on $\pi_* X$ and because L is level preserving, conjugation by L induces the identity on $\text{Wh}_2(\pi_1 X)$ and $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \widehat{\chi}$, so $\Sigma(\widetilde{J}) = \Sigma(\overline{J})$ and $\Theta(\widetilde{J}) = \Theta(\overline{J})$.

Further \widetilde{J} is isotopic to J in $\mathcal{P}(X)$ via the path

$$J_\theta = R_\theta^{-1} \circ J \circ R_\theta \in \mathcal{P}(X)$$

where $\theta \in [0, \pi]$. Note that $J_0 = J$ and $J_\pi = \widetilde{J}$. Hence applying Proposition 2.8.8 we have $\Sigma(J) = \Sigma(\widetilde{J}) = \Sigma(\overline{J}) = (-1)^n \overline{\Sigma(\overline{J})}$ and when $J \in \ker \Sigma$, $\widehat{\Theta}(J) = \widehat{\Theta}(\widetilde{J}) = \widehat{\Theta}(\overline{J}) = (-1)^n \overline{\Sigma(\overline{J})}$ as required. \square

2.8.3 Diffeomorphisms of $X^4 \times I$

In this section we prove Theorem F which gives diffeomorphisms of the 5-manifold $X \times I$, for X a 4-manifold, which are pseudo-isotopic but not isotopic for the identity. We do not use the results of this section elsewhere so a reader uninterested in 5-manifolds may skip this subsection entirely.

We recall the following result from Hatcher.

Proposition 2.8.10. [Hat73, Lemma 5.2] *Let X be an n -manifold, $n \geq 5$. Then*

$$B_n^2 \subset \Sigma(\mathcal{J}) \text{ and}$$

$$B_n^1 \subset \Theta(\mathcal{J} \cap \ker \Sigma).$$

Recall that we defined the subgroups B_n^i at the beginning of Section 2.8.2.

This uses the surjectivity of Σ and Θ . In 4-dimensions the following weaker statement still holds.

Proposition 2.8.11. *Let X be a 4-manifold. Then*

$$\begin{aligned} \{\theta + \bar{\theta} \mid \theta \in \Sigma(\mathcal{P})\} &\subset \Sigma(\mathcal{J}) \text{ and} \\ \{\theta + \bar{\theta} \mid \theta \in \widehat{\Theta}(\ker \Sigma)\} &\subset \widehat{\Theta}(\mathcal{J} \cap \ker \Sigma). \end{aligned}$$

Combining this with Theorem C gives the following corollary.

Corollary 2.8.12. *Let X be a 4-manifold. Then*

$$\{\theta + \bar{\theta} \mid \theta \in \Xi\} \subset \widehat{\Theta}(\mathcal{J} \cap \ker \Sigma)$$

where we consider Ξ in the quotient $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \widehat{\chi}$.

The proof for 4-dimensions is the same as that in high dimensions. Let X be an n -manifold, and let F be a pseudo-isotopy of X . Let $p_1: X \times I \rightarrow X \times I$ be the map which sends (x, s) to $(x, s/2)$, and let $p_2: X \times I \rightarrow X \times I$ be the map which sends (x, s) to $(x, 1 - s/2)$. We form the *double* of F , $2F \in \mathcal{J}(X)$ via

$$(2F)(x, s) = \begin{cases} p_1 \circ F(x, 2s) & s \leq 1/2 \\ p_2 \circ F(x, 2 - 2s) & s > 1/2 \end{cases}$$

Where p_X is the projection of $X \times I$ onto X . That is we compress F into the first half of the interval, and \bar{F} into the second half. It is clear that $2F \in \text{Diff}(X \times I, \partial(X \times I)) = \mathcal{J}(X)$.

Now as in [Hat73, Corollary 4.5] and [Hat73, Lemma 5.2] we have

$$\Sigma(2F) = \Sigma(F) + (-1)^n \overline{\Sigma(F)}$$

and

$$\widehat{\Theta}(2F) = \widehat{\Theta}(F) + (-1)^n \overline{\widehat{\Theta}(F)}.$$

Using $2F$ we can see the above elements of $\text{Wh}_2(\pi_1 X)$ and $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_1 X)$ in the image of \mathcal{J} under Σ and $\widehat{\Theta}$ as required.

We can now prove the an analogue of [Hat73, Corollary 4.5] for 5-manifolds.

Theorem F. *Suppose X is a 4-manifold which contains an element $\sigma \in \pi_2(X)$ with $w_2^X(\sigma) \neq 0$, and an element $\gamma \in \pi_1 X$ such that γ and γ^{-1} are not conjugate, and suppose also that either $k_1 X = 0$ or $K_3 \mathbb{Z}[\pi_1 X] = 0$. Then in $\text{Diff}(X \times I, \partial(X \times I))$ there exist diffeomorphisms pseudo-isotopic to the identity but not isotopic to it.*

Proof. By Corollary 2.6.4 identify $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$ with

$$(\mathbb{Z}_2 \times \pi_2 X)[\pi_1 X] / \langle \alpha\sigma - \alpha^T \tau \sigma \tau^{-1}, \beta \cdot 1 \mid \alpha, \beta \in \mathbb{Z}_2 \times \pi_2 X, \tau, \sigma \in \pi_1 X \rangle$$

and $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2)$ with

$$\mathbb{Z}_2[\pi_1 X] / \langle \sigma - \tau \sigma \tau^{-1}, 1 \mid \tau, \sigma \in \pi_1 X \rangle = \bigoplus_{\text{Conj}(\pi_1 X) \neq 1} \mathbb{Z}_2$$

where $\text{Conj}(\pi_1 X) \neq 1$ is the set of conjugacy classes of $\pi_1 X$ which are not the conjugacy class of 1.

By Theorem C there exists a pseudo-isotopy F of X , with $\Theta(F) = (1 + \sigma)\gamma$. Considering only $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2)$, this means F has $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2)$ invariant γ .

We proceed as in [Hat73, Corollary 4.5]. Consider the double

$$2F \in \text{Diff}(X \times I, \partial(X \times I)).$$

If we consider this as a pseudo-isotopy of X , then it has $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2)$ invariant $\gamma + \gamma^{-1}$.

Suppose $2F \in \text{Diff}(X \times I, \partial(X \times I))$ is isotopic to the identity. Then $2F$ would be isotopic as a pseudo-isotopy to the identity, so $\Theta(F) = 0$. But the $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2)$ invariant of $2F$ is $\gamma + \gamma^{-1}$, and by the assumption that γ and γ^{-1} are not conjugate, $\gamma + \gamma^{-1}$ is not conjugate to 1, so does not vanish in $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2)$, so this is a contradiction.

Hence it suffices to prove that $2F$, considered as a diffeomorphism of $X \times I$, is pseudo-isotopic to the identity. That is we must construct a pseudo-isotopy of $X \times I$ from the identity to $2F$, namely a diffeomorphism D of $(X \times I) \times I$ with $D|_{(X \times I) \times 1} = F$ and $(X \times I) \times 0$ the identity. For this construction, the suspension SF defined in [HW73, Chapter 1, Section 5] suffices. \square

2.8.4 Diffeomorphisms of $S^1 \times S^2 \times I$

We begin with the example $X = S^1 \times S^2 \times I$. We identify $\pi_2 X = \mathbb{Z}$ and $\pi_1 X$ with the multiplicative infinite cyclic group $\{t^n \mid n \in \mathbb{Z}\}$ so we may identify $(\pi_2 X)[\pi_1 X]$ with $\mathbb{Z}[t^\pm]$. We note that the action of $\pi_1 X$ on $\pi_2 X$ is trivial (one can see this in $S^1 \times S^2$).

Since by Section 2.1.4

$$k_1 X \in H^3(\pi_1 X; \pi_2 X) = H^3(\mathbb{Z}; \mathbb{Z}) = H^3(S^1; \mathbb{Z}) = 0$$

it follows that $k_1 X = 0$. Hence we can consider

$$\Theta(\ker \Sigma) \subset \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) = \text{Wh}_1(\pi_1 X; \pi_2 X) \oplus \text{Wh}_1(\pi_1 X; \mathbb{Z}_2).$$

By Proposition 2.6.4 we have

$$\begin{aligned} \text{Wh}_1(\pi_1 X; \pi_2 X) &= \mathbb{Z}[t] / \langle nt^a - n^{t^b} t^b t^a t^{-b}, n \cdot 1 \mid a, b, n \in \mathbb{Z} \rangle \\ &= \mathbb{Z}[t] / \langle n \cdot 1 \mid n \in \mathbb{Z} \rangle \\ &= \bigoplus_{i \in \mathbb{Z}^\times} \mathbb{Z} t^i. \end{aligned}$$

Similarly

$$\begin{aligned} \text{Wh}_1(\pi_1 X; \mathbb{Z}_2) &= \mathbb{Z}_2[t] / \langle mt^a - n^{t^b} t^b t^a t^{-b}, m \cdot 1 \mid a, b \in \mathbb{Z}, m \in \mathbb{Z}_2 \rangle \\ &= \mathbb{Z}_2[t] / \langle m \cdot 1 \mid m \in \mathbb{Z}_2 \rangle \\ &= \bigoplus_{i \in \mathbb{Z}^\times} \mathbb{Z}_2 t^i. \end{aligned}$$

Hence

$$\begin{aligned} \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) &= \text{Wh}_1(\pi_1 X; \mathbb{Z}_2) \oplus \text{Wh}_1(\pi_1 X; \pi_2 X) \\ &= \bigoplus_{i \in \mathbb{Z}^\times} \mathbb{Z}_2 t^i \oplus \bigoplus_{i \in \mathbb{Z}^\times} \mathbb{Z} t^i \\ &= \bigoplus_{i \in \mathbb{Z}^\times} (\mathbb{Z}_2 \times \mathbb{Z}) t^i. \end{aligned}$$

By Proposition 2.8.8 we also have that

$$\Theta(\mathcal{J}((S^1 \times S^2) \times I)) \subset \mathcal{Z}_4((S^1 \times S^2) \times I) = \{\theta \in \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 X) \mid \theta = \bar{\theta}\}.$$

For $(m, n) \cdot t^a \in \bigoplus_{i \in \mathbb{Z}^\times} (\mathbb{Z}_2 \times \mathbb{Z})t^i = \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$ note that

$$\overline{(m, n)t^a} = (m + w_2^X(m), -w_1^X(m)n^{t^{-a}})t^{-a} = (m, -n)t^{-a}$$

since w_1^X and w_2^X are trivial for $S^1 \times S^2 \times I$.

Since $a \neq 0$ we never have that $(m, n)t^a = (m, -n)t^{-a}$ so it is clear that

$$Z_4((S^1 \times S^2) \times I) = \{b \in \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) \mid b = \bar{b}\} = \langle (m, n)t^a + (m, -n)t^{-a} \rangle.$$

Hence quotienting out by $Z_4(S^1 \times S^2 \times I)$ just identifies $\mathbb{Z}_2 \times \mathbb{Z}t^a$ with $\mathbb{Z}_2 \times \mathbb{Z}t^{-a}$, so we have a map

$$\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \Theta(\mathcal{J} \cap \ker \Sigma) \xrightarrow{q} \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / Z_4(X) = \bigoplus_{i \in \mathbb{Z}_{>0}} (\mathbb{Z}_2 \times \mathbb{Z})t^i.$$

By Corollary C we have $F \in \ker \Sigma \subset \mathcal{P}(X)$ with $\Theta(F) = (0, n)t^a$, and so it follows we have $f \in \ker \Sigma \subset \text{Diff}_{PI}(S^1 \times S^2 \times I, \partial(S^1 \times S^2 \times I))$ with $\Theta(f) = (0, n)t^a$. Let p_2 be the projection $p_2: \bigoplus_{i \in \mathbb{Z}_{>0}} (\mathbb{Z}_2 \times \mathbb{Z})t^i \rightarrow \bigoplus_{i \in \mathbb{N}} \mathbb{Z}$. The $S^1 \times S^2 \times I$ case of Theorem A follows letting $K = \ker \Sigma$, and $\Theta' = p_2 \circ q \circ \Theta$.

Theorem A. *Let X be either the 4-manifold $S^1 \times S^2 \times I$ or $(M_1 \# M_2) \times I$, for M_1, M_2 closed, orientable, aspherical 3-manifolds. Then there is a subgroup $K \leq \pi_0 \text{Diff}_{PI}(X, \partial X)$ and a surjective map*

$$\Theta': K \longrightarrow \bigoplus_{i \in \mathbb{N}} \mathbb{Z}.$$

Hence there are infinitely many distinct isotopy classes of diffeomorphisms of X fixing the boundary, which are pseudo-isotopic to the identity.

Note that as $w_2^X = 0$ for $S^1 \times S^2 \times I$ we do not know how to realise the \mathbb{Z}_2 part of $\bigoplus_{i \in \mathbb{Z}_{>0}} (\mathbb{Z}_2 \times \mathbb{Z})t^i$ which arises from the framing.

2.8.5 Diffeomorphisms of the connect sum of aspherical 3-manifolds times I

In this section we produce diffeomorphisms for $X = (M_1 \# M_2) \times I$, where M_i are closed, orientable, aspherical 3-manifolds. The condition of being aspherical is equivalent to being

irreducible with an infinite fundamental group; this follows from the sphere theorem, see [Hem76, Theorem 4.3]. Many examples of such 3-manifolds exist, including $\Sigma_g \times S^1$ for Σ_g a surface of genus $g > 0$, as well as many hyperbolic 3-manifolds.

Note that aspherical 3-manifolds M_i have torsion free fundamental group. To see this note that M_i is $K(\pi_1 M_i, 1)$ space. If $G \leq \pi_1 M_i$ is a cyclic subgroup, let \widetilde{M}_i be the corresponding cover of M_i . Then \widetilde{X} is a $K(G, 1)$ space so $H_i(G, \mathbb{Z}) = H_i(\widetilde{M}_i, \mathbb{Z}) = 0$ for $i > 3$, which is only possible if G is infinite; see [Hat02, Proposition 2.45].

We first compute $\pi_1 X$ and $\pi_2 X$ along with the action. Let $M = M_1 \# M_2$. It is clear that $\pi_i X = \pi_i M$ for all i . It is also clear that $\pi_1 M = \pi_1 M_1 * \pi_1 M_2$. To compute $\pi_2 M$ we consider the universal cover $p: \widetilde{M} \rightarrow M$. Writing $M = (M_1 \setminus B^3) \cup_{S^2} (M_2 \setminus B^3)$ we denote $Y_i = p^{-1}(M_i \setminus B^3)$. Considering the action of $\pi_1 M$ on \widetilde{M} we can make the following identifications

$$\begin{aligned} Y_1 &= \bigsqcup_{\pi_1 M / \pi_1 M_1} (\widetilde{M_1 \setminus B^3}), \\ Y_2 &= \bigsqcup_{\pi_1 M / \pi_1 M_2} (\widetilde{M_2 \setminus B^3}), \\ Y_1 \cap Y_2 &= \bigsqcup_{\pi_1 M} S^2. \end{aligned}$$

We write the Mayer-Vietoris sequence for $\widetilde{M} = Y_1 \cup Y_2$ with coefficients in \mathbb{Z}

$$0 = H_3(\widetilde{M}) \rightarrow H_2(Y_1 \cap Y_2) \xrightarrow{(j_1, -j_2)} H_2(Y_1) \oplus H_2(Y_2) \xrightarrow{i_1 + i_2} H_2(\widetilde{M}) \rightarrow H_1(Y_1 \cap Y_2) = 0.$$

Note that $H_3(\widetilde{M}) = 0$ as $\pi_1 M$ is infinite, so \widetilde{M} is non compact. Using the identifications above we obtain

$$\begin{aligned} 0 &\rightarrow \mathbb{Z}[\pi_1 M] \xrightarrow{(j_1, -j_2)} \left(\mathbb{Z}[\pi_1 M] \otimes_{\mathbb{Z}\pi_1 M_1} H_2(\widetilde{M_1 \setminus B^3}) \right) \oplus \left(\mathbb{Z}\pi_1 M \otimes_{\mathbb{Z}\pi_1 M_2} H_2(\widetilde{M_2 \setminus B^3}) \right) \\ &\rightarrow H_2(\widetilde{M}) \rightarrow 0. \end{aligned}$$

To calculate $H_2(\widetilde{M_i \setminus B^3})$, note that since M_i is aspherical, $H_2(\widetilde{M_i}) = \pi_2 M_i = 0$. Note also that

$$\widetilde{M_1 \setminus B^3} = \widetilde{M_1} \setminus \bigcup_{g \in \pi_1 M} g \widetilde{B^3}$$

for $\widetilde{B^3}$ some lift of the ball B^3 (which we note is also a ball). Hence $H_2(\widetilde{M_i \setminus B^3})$ is generated by the boundaries of these balls, and the only possible relation between these generators comes from taking the boundary of a 3-chain corresponding to the entire 3-manifold $\widetilde{M_i \setminus B^3}$. There is only such a 3-chain if $\widetilde{M_i \setminus B^3}$ is compact, which does not hold since $\pi_1 M_i$ is infinite. Hence $H_2(\widetilde{M_i \setminus B^3}) = \mathbb{Z}[\pi_1 M_i]$. Substituting this into the short exact sequence we obtain:

$$0 \rightarrow \mathbb{Z}[\pi_1 M] \xrightarrow{(1, -1)} \mathbb{Z}[\pi_1 M] \oplus \mathbb{Z}[\pi_1 M] \xrightarrow{i_1 + i_2} H_2(\widetilde{M}) \rightarrow 0.$$

Hence,

$$\pi_2(M_1 \# M_2) = H_2(\widetilde{M_1 \# M_2}) = \mathbb{Z}[\pi_1 M] = \mathbb{Z}[\pi_1 M_1 * \pi_1 M_2]$$

with the action of $\pi_1 M$ on $\pi_2 M$ given by the obvious left multiplication.

By Proposition 2.6.4 we have

$$\begin{aligned} \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) &= (\mathbb{Z}_2 \times \mathbb{Z}[\pi_1 M])[\pi_1 M] / \langle (m, ng)a - (m, (ng)^b)bab^{-1}, (m, ng)1 \rangle \\ &= (\mathbb{Z}_2 \times \mathbb{Z}[\pi_1 M])[\pi_1 M] / \langle (m, ng)a - (m, nbg)bab^{-1}, (m, ng)1 \rangle. \end{aligned}$$

Identifying $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$ with this quotient of $(\mathbb{Z}_2 \times \mathbb{Z}[\pi_1 M])[\pi_1 M]$ consider the surjective map,

$$\begin{aligned} q: \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) &\longrightarrow \bigoplus_{S \in \text{Conj}(\pi_1 X) \neq 1} (\mathbb{Z}_2 \times \mathbb{Z})S \\ q: (m, ng)a &\longmapsto \begin{cases} (m, n) \text{Cl}(a) & \text{if } \text{Cl}(a) \neq 1 \\ 0 & \text{if } \text{Cl}(a) = 1. \end{cases} \end{aligned}$$

where $\text{Conj}(\pi_1 X) \neq 1$ denotes the set of conjugacy classes which are not the conjugacy class of 1, and $\text{Cl}(a)$ denotes the conjugacy class of a . To see this is well defined we note that it vanishes on both relations since

$$\begin{aligned} (m, ng)a - (m, nbg)bab^{-1} &\longmapsto (m, n) \text{Cl}(a) - (m, n) \text{Cl}(bab^{-1}) \\ &= (m, n) \text{Cl}(a) - (m, n) \text{Cl}(a) = 0. \end{aligned}$$

Since $\pi_1 M = \pi_1 M_1 * \pi_1 M_2$, and $\pi_1 M_i$ are infinite there are many conjugacy classes in $\pi_1 M$.

Ultimately we wish to consider the quotient

$$(\mathrm{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \widehat{\chi}) / \widehat{\Theta}(\mathcal{J}(X) \cap \ker \Sigma).$$

Note by Proposition 2.8.9 we have

$$\widehat{\Theta}(\mathcal{J}(X) \cap \ker \Sigma) \leq Z_4^1(X) = \{\theta \in \mathrm{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \widehat{\chi} \mid \theta = \bar{\theta}\}.$$

Let

$$\check{Z} = \{\theta \in \mathrm{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) \mid \theta = \bar{\theta}\} \leq \mathrm{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$$

and let r be the quotient map $r: \mathrm{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) \rightarrow \mathrm{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \check{Z}$. Then it is clear that

$$(\mathrm{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \widehat{\chi}) / Z_4^1(X) = (\mathrm{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \check{Z}) / r(\widehat{\chi}).$$

Our approach for the remainder of this section is to use the map

$$(\mathrm{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \widehat{\chi}) / \widehat{\Theta}(\mathcal{J}(X) \cap \ker \Sigma) \rightarrow (\mathrm{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \check{Z}) / r(\widehat{\chi})$$

to understand $(\mathrm{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \widehat{\chi}) / \widehat{\Theta}(\mathcal{J}(X) \cap \ker \Sigma)$.

Again identifying $\mathrm{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X)$ with a quotient of $(\mathbb{Z}_2 \times \mathbb{Z}[\pi_1 X]) / [\pi_1 X]$, first note that

$$\overline{(m, ng)a} = (m + w_2^X(ng), -w_1^X(a)(ng)^{a^{-1}})a^{-1} = (m, -na^{-1}g)a^{-1}$$

since w_1^X and w_2^X are trivial in M as it is an orientable 3-manifold, and so w_1^X and w_2^X are also trivial in $X = M \times I$. Noting that

$$q(\overline{(m, ng)a}) = (m, -n) \mathrm{Cl}(a^{-1}),$$

we define an involution on $\bigoplus_{S \in \mathrm{Conj}(\pi_1 X) \neq 1} (\mathbb{Z}_2 \times \mathbb{Z})S$ via

$$\begin{aligned} \bar{\cdot}: \bigoplus_{S \in \mathrm{Conj}(\pi_1 X) \neq 1} (\mathbb{Z}_2 \times \mathbb{Z})S &\rightarrow \bigoplus_{S \in \mathrm{Conj}(\pi_1 X) \neq 1} (\mathbb{Z}_2 \times \mathbb{Z})S \\ \bar{\cdot}: (m, n) \mathrm{Cl}(a) &\mapsto (m, -n) \mathrm{Cl}(a^{-1}). \end{aligned}$$

Now $q(\bar{a}) = \overline{q(a)}$, so $q(\check{Z}) = \{s \in \bigoplus_{S \in \mathrm{Conj}(\pi_1 X) \neq 1} (\mathbb{Z}_2 \times \mathbb{Z})S \mid s = \bar{s}\}$. Given a conjugacy class $S \in \mathrm{Conj}(\pi_1 X)$ we denote $\bar{S} = \mathrm{Cl}(a^{-1})$ where $S = \mathrm{Cl}(a)$.

We claim that

$$q(\tilde{Z}) = \langle (m, n)S + (m, -n)\bar{S}, (m, 0)P \mid m \in \mathbb{Z}_2, n \in \mathbb{Z}, S, P \in \text{Conj}(\pi_1 X), \bar{P} = P \rangle$$

Note that if $\text{Cl}(a^{-1}) = \text{Cl}(b^{-1})$ then $a^{-1} = rb^{-1}r^{-1}$, so $a = rbr^{-1}$ and $\text{Cl}(a) = \text{Cl}(b)$, so $\bar{S} = \bar{P} \implies S = P$. To prove this claim, we can write any element in

$$\bigoplus_{S \in \text{Conj}(\pi_1 X)^{\neq 1}} (\mathbb{Z}_2 \times \mathbb{Z})S$$

as $\sum_i (m_i, n_i)S_i$, where the S_i are distinct. If

$$\sum_i (m_i, n_i)S_i = \overline{\sum_i (m_i, n_i)S_i} = \sum_i (m_i, -n_i)\bar{S}_i$$

then since $\bar{\cdot}$ is injective, there is a permutation σ which pairs S_i with the unique $S_{\sigma(i)}$ such that $S_i = \overline{S_{\sigma(i)}}$. We hence also see that $m_i = m_{\sigma(i)}$ and $n_i = -n_{\sigma(i)}$; note that if $i = \sigma(i)$ then $n_i = -n_i = 0$. Hence we can rewrite the sum as

$$\begin{aligned} \sum_i (m_i, n_i)S_i &= \sum_{i, i=\sigma(i)} (m_i, 0)S_i + \sum_{i, i<\sigma(i)} (m_k, n_k)S_k + (m_{\sigma(i)}, n_{\sigma(i)})S_{\sigma(i)} \\ &= \sum_{i, i=\sigma(i)} (m_i, 0)S_i + \sum_{i, i<\sigma(i)} (m_i, n_i)S_i + (m_i, -n_i)\bar{S}_i \end{aligned}$$

which is the sum of generators of the required form, proving our claim.

Hence we can see that quotienting $\bigoplus_{S \in \text{Conj}(\pi_1 X)^{\neq 1}} (\mathbb{Z}_2 \times \mathbb{Z})S$ by $p(\tilde{Z})$ identifies $(\mathbb{Z}_2 \times \mathbb{Z})S$ with $(\mathbb{Z}_2 \times \mathbb{Z})\bar{S}$ when $S \neq \bar{S}$, and kills the \mathbb{Z}_2 part when $S = \bar{S}$, that is

$$\left(\bigoplus_{S \in \text{Conj}(\pi_1 X)^{\neq 1}} (\mathbb{Z}_2 \times \mathbb{Z})S \right) / p(\tilde{Z}) = \bigoplus_{\substack{S \in \text{Conj}(\pi_1 X)^{\neq 1}, \\ S = \bar{S}}} \mathbb{Z}S \oplus \bigoplus_{\substack{[S] \in \text{Conj}(\pi_1 X)^{\neq 1} / \sim, \\ S \neq \bar{S}}} (\mathbb{Z}_2 \times \mathbb{Z})S$$

where \sim is the equivalence relation on $\text{Conj}(\pi_1 X)^{\neq 1}$ given by $S \sim \bar{S}$. Clearly q induces a surjective map between the quotients

$$\tilde{q}: \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \tilde{Z} \rightarrow \bigoplus_{\substack{S \in \text{Conj}(\pi_1 X)^{\neq 1}, \\ S = \bar{S}}} \mathbb{Z}S \oplus \bigoplus_{\substack{[S] \in \text{Conj}(\pi_1 X)^{\neq 1} / \sim, \\ S \neq \bar{S}}} (\mathbb{Z}_2 \times \mathbb{Z})S.$$

As previously, we will see that we are able to realise the $0 \times \mathbb{Z}$ part of the second summand, as well as the \mathbb{Z} part of the first summand, however, it will be easier to come up with conjugacy classes with $S \neq \bar{S}$.

Conjugacy classes in free products.

In order to come up with suitable conjugacy classes, in this subsection we prove the following

Proposition 2.8.13. *For all $n \in \mathbb{N}$ there exists $a_n \in \pi_1 X = \pi_1 M_1 * \pi_1 M_2$ such that a_n and a_m are not conjugate for $n \neq m$, and that a_n and a_m^{-1} are not conjugate $\forall n, m \in \mathbb{N}$; in particular a_n is not conjugate to its inverse. Hence there are infinitely many distinct equivalence classes of conjugacy classes $[S] \in \text{Conj}(\pi_1 X) / \sim$ such that $S \neq \bar{S}$.*

This gives us the following corollary.

Corollary 2.8.14. *The abelian group*

$$\bigoplus_{S \in \text{Conj}(\pi_1 X)^{\neq 1}, S = \bar{S}} \mathbb{Z}S \oplus \bigoplus_{[S] \in \text{Conj}(\pi_1 X)^{\neq 1} / \sim, S \neq \bar{S}} (\mathbb{Z}_2 \times \mathbb{Z})S$$

has infinite rank, hence $\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \check{Z}$ also has infinite rank as \tilde{q} is surjective.

Proposition 2.8.13 follows easily from the following fact about free products of groups.

Lemma 2.8.15. *Given groups A and B , let $a \in A, b \in B$ with $a, b \neq 1$, then $(ab)^n \in A * B$ are distinct conjugacy classes in $A * B$ for all $n \in \mathbb{N}$. If additionally one of a or b is not of order two, then $(ab)^n \in A * B$ are distinct conjugacy classes in $A * B$ for all $n \in \mathbb{Z}$.*

Taking $A = \pi_1 M_1, B \in \pi_1 M_2$, setting $a_n = (ab)^n$ proves Proposition 2.8.13; note that $\pi_1 M_i$ is torsion free so has no order two elements. We give a proof of Lemma 2.8.15 below.

Proof of Lemma 2.8.15. We define a length function $l : \text{Conj}(A * B)^{\neq 1} \rightarrow \mathbb{N}$ by defining $l(S)$ to be the minimum n such that we can write $S = \text{Cl}(c_1 c_2 \cdots c_n)$ such that $1 \neq c_i \in A$ for i odd and $1 \neq c_i \in B$ for i even, or $1 \neq c_i \in A$ for i even and $1 \neq c_i \in B$ for i odd. That is, it is the minimum n such that we can give as an alternating product in elements of A and B .

Note that in the free product $A * B$ it is clear that if $c_1 \dots c_n = d_1 \dots d_n$ for both c_i and d_i alternating in A and B , then $c_i = d_i \forall i$.

For $a_i \in A, b_i \in B$, we claim that $l(\text{Cl}(a_1 b_1 a_2 b_2 \cdots a_n b_n))$ is $2n$. To prove this, suppose

$$r a_1 b_1 a_2 b_2 \dots a_n b_n r^{-1} = c_1 c_2 \dots c_m$$

for $m < 2n$ and $c_1 c_2 \dots c_m$ an alternating sum in A and B . We can also write $r = r_1 r_2 \dots r_k$ as an alternating sum in A and B , so we have

$$r_1 r_2 \dots r_k a_1 b_1 a_2 b_2 \dots a_n b_n r_k^{-1} r_{k-1}^{-1} \dots r_1^{-1} = c_1 c_2 \dots c_m.$$

Suppose first that $r_n \in B$. Then $r_1 \dots r_n a_1 b_1 a_2 b_2 \dots a_n b_n$ is alternating. If $b_n r_k^{-1} \neq 1$ then

$$r_1 r_2 \dots r_k a_1 b_1 a_2 b_2 \dots a_n (b_n r_k^{-1}) r_{k-1}^{-1} \dots r_1^{-1}$$

is alternating of length $2n + 2k - 1$, so $m = n + 2k - 1 > n$ which is a contradiction. Hence $b_n r_k^{-1} = 1$. Cancelling $b_n r_k^{-1} = 1$ and repeating this argument on $r_k a_1 b_1 a_2 b_2 \dots a_n r_{k-1}^{-1} \dots r_1^{-1}$ we prove that $a_n r_{k-1}^{-1} = 1$.

When $k \leq 2n$, we can repeat this argument to prove that $r_k^{-1} r_{k-1}^{-1} \dots r_1^{-1}$ cancels with the right k terms of $a_1 b_1 \dots a_n b_n$ at which point the remaining product $r_1 \dots r_n a_1 b_2 \dots$ is alternating and $m = 2n$ which is a contradiction.

When $k > 2n$ we cancel every term in $a_1 b_1 \dots a_n b_n$ and are left with

$$r_1 r_2 \dots r_k r_{k-2n}^{-1} \dots r_1^{-1}.$$

Now if $r_k r_{k-2n}^{-1} \neq 1$, then $r_1 r_2 \dots (r_k r_{k-2n}^{-1}) r_{k-2n-1}^{-1} \dots r_1^{-1}$ is again alternating and has length $k + k - 2n - 1 \geq n$, so $m = k + k - 2n - 1$ which is a contradiction. Hence we must have $r_k r_{k-2n}^{-1} = 1$.

Repeating this argument we see that all $k - 2n$ terms of $r_{k-2n}^{-1} \dots r_1^{-1}$ must cancel the right $k - 2n$ terms of $r_1 r_2 \dots r_k$ and are left with $r_1 \dots r_{2n}$ which is alternating of length $2n$ so again a contradiction.

If we suppose instead that $r_n \in A$, we see that $a_1 b_1 a_2 b_2 \dots a_n b_n r_k^{-1} r_{k-1}^{-1} \dots r_1^{-1}$ is alternating, and run the same argument proving that $r_k a_1 = 1$ and so on, again arriving at a contradiction.

Hence we cannot write $r a_1 b_1 a_2 b_2 \dots a_n b_n r^{-1} = c_1 c_2 \dots c_m$ for $m \leq n$. This completes the proof of our claim that $l(\text{Cl}(a_1 b_1 a_2 b_2 \dots a_n b_n))$ is $2n$.

Now since $l(\text{Cl}((ab)^n)) = 2n$, we have that $(ab)^n$ are in distinct conjugacy classes for all $n \in \mathbb{N}$ as required.

Suppose now that one of a or b is not order two. To prove that $(ab)^n$ are in distinct conjugacy classes for all $n \in \mathbb{Z}$ it is sufficient to prove that $(ab)^n$ is not conjugate to $(ab)^{-n}$.

Suppose that $r(ab)^n r^{-1} = (ab)^{-n} = (b^{-1}a^{-1})^n$ writing $r = r_1 r_2 \dots r_k$ as an alternating product in A and B , we have

$$r_1 r_2 \dots r_k (ab)^n r_k^{-1} \dots r_2^{-1} r_1^{-1} = (b^{-1}a^{-1})^n$$

Suppose $r_k \in B$. Then we must have that $r_1 \in B$ as the right hand term starts with an element of B . Hence k is odd. Now the only way the alternating length of the left can agree with the alternating length of the right is if $r_k^{-1} \dots r_2^{-1} r_1^{-1}$ cancels with the right hand k terms of $(ab)^n$, hence $r_k^{-1} r_{k-1}^{-1} \dots r_1 = b^{-1} a^{-1} b^{-1} \dots a^{-1} b^{-1}$. Now we have

$$r_k (ab)^n r_k^{-1} \dots r_2^{-1} r_1^{-1} = (ba)^n = (b^{-1}a^{-1})^n$$

which is only possible if $a = a^{-1}$ and $b = b^{-1}$ so a and b are order two which is a contradiction.

Suppose instead that $r_k \in A$. Then we see that $r_1 \in A$ as the right hand term ends with a term of A . Hence again k is odd. Similarly to the previous argument we see that r must cancel with the left hand k terms and again conclude that

$$r_k (ab)^n r_k^{-1} \dots r_2^{-1} r_1^{-1} = (ba)^n$$

again leading to a contradiction. □

The rank of $K_3\mathbb{Z}[\pi_1 M_1 \# M_2]$

To complete our argument, we must quotient out by $r(\widehat{\chi}) = r(\chi(K_3\mathbb{Z}[\pi_1 X]) + \overline{\chi(K_3\mathbb{Z}[\pi_1 X])})$.

We will prove that $K_3\mathbb{Z}[\pi_1 X]$ has finite rank, and so $r(\widehat{\chi})$ also has finite rank.

Proposition 2.8.16. *If $M = M_1 \# M_2$ is a 3-manifold that is the connect sum of two aspherical 3-manifolds then $K_3\mathbb{Z}[\pi_1 M]$ has rank two.*

Proof. In order to compute $K_3\mathbb{Z}[\pi_1 M]$, we first note that the Farrell Jones conjecture holds for 3-manifold groups; see [BFL14, Corollary 0.3]. Since $\pi_1 M = \pi_1 M_1 * \pi_1 M_2$ is torsion free there is an isomorphism

$$H_n(B(\pi_1 M); \mathbf{K}(\mathbb{Z})) \xrightarrow{\cong} K_n\mathbb{Z}[\pi_1 M]$$

where $\mathbf{K}(\mathbb{Z})$ is the K -Theory spectrum of \mathbb{Z} , and $H_n(-; \mathbf{K}(\mathbb{Z}))$ is the generalised homology

theory associated to this spectrum; see [LR05] for further details on this, and the Farrell Jones conjecture.

We proceed to calculate $H_3(B(\pi_1 M); \mathbf{K}(\mathbb{Z}))$. We first further simplify things by noting that since M_i are aspherical we have that $B(\pi_1 M_i) = M_i$ and so

$$B(\pi_1 M) = B(\pi_1 M_1 * \pi_1 M_2) = B(\pi_1 M_1) \vee B(\pi_1 M_2) = M_1 \vee M_2.$$

Using the axioms for generalised homology theories we have that

$$H_n(B(\pi_1 M); \mathbf{K}(\mathbb{Z})) \cong H_n(M_1; \mathbf{K}(\mathbb{Z})) \oplus H_n(M_2; \mathbf{K}(\mathbb{Z}))$$

hence it will be sufficient to calculate $H_n(M_i; \mathbf{K}(\mathbb{Z}))$.

For any generalised homology theory there is an Atiyah-Hirzebruch spectral sequence with E_2 page given by

$$E_2^{p,q} = H_p(M_i; K_q(\mathbb{Z}))$$

where $H_p(M_i; K_p(\mathbb{Z}))$ is usual singular homology with coefficients, and $K_p(\mathbb{Z})$ are the algebraic K-theory groups of \mathbb{Z} ; see for example [DK01] for details on the Atiyah-Hirzebruch spectral sequence. This spectral sequence converges to $H_n(B(\pi_1 X); \mathbf{K}(\mathbb{Z}))$ in the following sense; on each homology group there is a filtration

$$H_n(M_i; \mathbf{K}(\mathbb{Z})) = \mathcal{F}_0^n \supset \mathcal{F}_1^n \supset \dots \supset \mathcal{F}_k^n = 0$$

and the E_∞ page of the spectral sequence gives

$$E_\infty^{n-j,j} = \mathcal{F}_j^n / \mathcal{F}_{j+1}^n.$$

We will use the first few terms of $K_q(\mathbb{Z})$, namely $K_0(\mathbb{Z}) = \mathbb{Z}$, $K_1(\mathbb{Z}) = \mathbb{Z}_2$, $K_2(\mathbb{Z}) = \mathbb{Z}_2$, $K_3(\mathbb{Z}) = \mathbb{Z}_{48}$; see for example [Wei05].

Below we write out the second page of the spectral sequence; in red we highlight those terms with $p + q = 3$, that will contribute to $H_3(M_i; \mathbf{K}(\mathbb{Z}))$, and we draw arrows where there is a non zero differential; note that for $p < 0$ and $p > 3$, $E_2^{p,q} = H_p(M_i; K_q(\mathbb{Z})) = 0$.

4	\vdots	\vdots	\vdots	\vdots
3	\mathbb{Z}_{48}	$H_1(M_i; \mathbb{Z}_{48})$	$H_2(M_i; \mathbb{Z}_{48})$	\mathbb{Z}_{48}
2	\mathbb{Z}_2	$H_1(M_i; \mathbb{Z}_2)$	$H_2(M_i; \mathbb{Z}_2)$	\mathbb{Z}_2
1	\mathbb{Z}_2	$H_1(M_i; \mathbb{Z}_2)$	$H_2(M_i; \mathbb{Z}_2)$	\mathbb{Z}_2
0	\mathbb{Z}	$H_1(M_i; \mathbb{Z})$	$H_2(M_i; \mathbb{Z})$	\mathbb{Z}
	0	1	2	3

We note that above the first row all groups are torsion, and indeed finite! The only group which is non torsion and which contributes to $H_3(M_i; \mathbf{K}(\mathbb{Z}))$ is $E_2^{3,0} = \mathbb{Z}$. Since the differential $d_2^{3,0}: \mathbb{Z} \rightarrow E_2^{1,1} = H_1(M_i; \mathbb{Z}_2)$, drawn in green, is a map into a finite group, we see that we must have $\ker(d_2^{3,0}) = \mathbb{Z}$, so $E_3^{3,0} = \mathbb{Z}$.

Now consider $d_3^{3,0}: \mathbb{Z} \rightarrow E_3^{0,2}$. since $E_3^{0,2}$ is a quotient of \mathbb{Z}_2 , it is either \mathbb{Z}_2 or 0, so again we see that $\ker(d_3^{3,0}) = \mathbb{Z}$ and that $E_4^{3,0} = \mathbb{Z}$. Since all differentials $d_n^{p,q}$ are zero for $n \geq 4$, $E_\infty^{p,q} = E_4^{p,q}$.

Hence $\mathcal{F}_0^3/\mathcal{F}_1^3 = \mathbb{Z}$, while $\mathcal{F}_1^3/\mathcal{F}_2^3 = E_4^{2,1}$, $\mathcal{F}_2^3/\mathcal{F}_3^3 = E_4^{1,2}$, $\mathcal{F}_3^3/\mathcal{F}_4^3 = E_4^{0,3}$ are all torsion and finite. We also have $\mathcal{F}_4^3/\mathcal{F}_5^3 = E_4^{-1,4} = 0$; since the filtration terminates with 0, it must be that $\mathcal{F}_4^3 = 0$. Hence \mathcal{F}_3^3 is finite, hence \mathcal{F}_2^3 is finite, and hence \mathcal{F}_1^3 is finite.

We have a short exact sequence

$$0 \rightarrow \mathcal{F}_1^3 \rightarrow \mathcal{F}_0^3 = H_3(M_i; \mathbf{K}(\mathbb{Z})) \rightarrow \mathcal{F}_0^3/\mathcal{F}_1^3 = \mathbb{Z} \rightarrow 0.$$

as all the groups are abelian and $\mathcal{F}_0^3/\mathcal{F}_1^3 = \mathbb{Z}$ is free abelian, the sequence splits and so

$$H_3(M_i; \mathbf{K}(\mathbb{Z})) = \mathbb{Z} \oplus \mathcal{F}_1^3.$$

Since \mathcal{F}_1^3 is finite, it follows that $H_3(M_i; \mathbf{K}(\mathbb{Z}))$ is rank one, and so $H_n(B\pi_1 M; \mathbf{K}(\mathbb{Z}))$ is rank two. □

Diffeomorphisms of $\text{Diff}_{PI}((M_1 \# M_2) \times I, \partial((M_1 \# M_2) \times I))$

Recall the map

$$\tilde{q}: \text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \check{Z} \rightarrow \bigoplus_{\substack{S \in \text{Conj}(\pi_1 X)^{\neq 1}, \\ S = \bar{S}}} \mathbb{Z}S \oplus \bigoplus_{\substack{[S] \in \text{Conj}(\pi_1 X)^{\neq 1} / \sim, \\ S \neq \bar{S}}} (\mathbb{Z}_2 \times \mathbb{Z})S.$$

Denote the target of this map by R . We can induce a map on the quotients

$$\tilde{q}: (\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \check{Z}) / r(\hat{\chi}) \rightarrow R / \tilde{q}(r(\hat{\chi}))$$

Recall that $(\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \check{Z}) / r(\hat{\chi}) = (\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \hat{\chi}) / Z_4^1(X)$ and that we have a map

$$(\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \hat{\chi}) / \hat{\Theta}(\mathcal{J}(X) \cap \ker \Sigma) \rightarrow (\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \hat{\chi}) / Z_4^1(X).$$

Composing this with \tilde{q} we obtain a map

$$q': (\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \hat{\chi}) / \hat{\Theta}(\mathcal{J}(X) \cap \ker \Sigma) \rightarrow R / \tilde{q}(r(\hat{\chi})).$$

The left hand side is precisely the target of

$$\hat{\Theta}: \ker \Sigma \subset \text{Diff}_{PI}(X, \partial X) \rightarrow (\text{Wh}_1(\pi_1 X; \mathbb{Z}_2 \times \pi_2 X) / \hat{\chi}) / \hat{\Theta}(\mathcal{J}(X) \cap \ker \Sigma).$$

Taking the composition $q' \circ \hat{\Theta}$ gives $q' \circ \hat{\Theta}: \ker \Sigma \rightarrow R / \tilde{q}(r(\hat{\chi}))$. By Theorem C the image of $q' \circ \hat{\Theta}$ contains

$$\left(\bigoplus_{\substack{S \in \text{Conj}(\pi_1 X)^{\neq 1}, \\ S = \bar{S}}} \mathbb{Z}S \oplus \bigoplus_{\substack{[S] \in \text{Conj}(\pi_1 X)^{\neq 1} / \sim, \\ S \neq \bar{S}}} (0 \times \mathbb{Z})S \right) / \tilde{q}(r(\hat{\chi})) \leq R / \tilde{q}(r(\hat{\chi})).$$

By Proposition 2.8.13

$$S = \bigoplus_{\substack{S \in \text{Conj}(\pi_1 X)^{\neq 1}, \\ S = \bar{S}}} \mathbb{Z}S \oplus \bigoplus_{\substack{[S] \in \text{Conj}(\pi_1 X)^{\neq 1} / \sim, \\ S \neq \bar{S}}} (0 \times \mathbb{Z})S$$

has infinite rank. By Proposition 2.8.16, $\hat{\chi} = \chi(K_3 \mathbb{Z}[\pi_1 M]) + \overline{\chi(K_3 \mathbb{Z}[\pi_1 M])}$ is at most rank four, and so $\tilde{q}(r(\hat{\chi}))$ is at most rank four. Hence $S / \tilde{q}(r(\hat{\chi}))$ has infinite rank, and so contains a subgroup isomorphic to $\bigoplus_{\mathbb{N}} \mathbb{Z}$. Letting $K = \ker \Sigma$ and letting Θ' be the composition of $q' \circ \Theta$ and projection onto the subgroup isomorphic to $\bigoplus_{\mathbb{N}} \mathbb{Z}$ yields the $(M_1 \# M_2) \times I$ case of Theorem A.

Chapter 3

Distances between surfaces in 4-manifolds

In Chapter 2 we dealt extensively with homotopies of surfaces in 4-manifolds. In this chapter we will expand on this topic further, our key result being Theorem G.

Theorem G. *If $\Sigma, \Sigma' \subset X$ are connected, smooth, properly embedded, oriented surfaces of the same genus then*

$$d_{\text{st}}(\Sigma, \Sigma') \leq d_{\text{sing}}(\Sigma, \Sigma') + 1.$$

See Section 1.2 for the definitions of the above distances.

Throughout X will be a smooth compact orientable 4-manifold possibly with boundary. We will consider immersed surfaces $\Sigma \subset X$ of genus g ; that is, Σ is the image of some generic immersion $f: S \rightarrow X$ where S is an abstract surface of genus g .

Recall the definitions of finger moves, Whitney moves, Whitney arcs, and Whitney framings from Section 2.1.2. Note that in Chapter 2 it was important that our finger move arcs be oriented, in this chapter we drop this requirement; as in Remark 2.1.4, reversing the orientation of the finger move arc does not change the resulting surface (up to homotopy of the surface), nor does it change the homotopy (up to isotopy of the homotopy).

Recall also that embedded surfaces $\Sigma, \Sigma' \subset X$ are regularly homotopic if there exists a smooth map $H: S \times [0, 1] \rightarrow X$ (where S is an abstract surface of genus g) where $H(-, 0)$ and $H(-, 1)$ are embeddings with $H(S, 0) = \Sigma$, $H(S, 1) = \Sigma'$, and $H(-, t)$ is an immersion

for all $t \in [0, 1]$. In the case X and S have boundary, we require the embeddings be proper embeddings, that $H(t, \partial S) \subset \partial X$ for all t , and that $H(t, -)$ is an embedding when restricted to some neighbourhood of ∂X for all t .

Remark 3.0.1. By theorems of Smale [Sma57] and Hirsch [Hir59] two embeddings of an orientable surface in an orientable 4-manifold are regularly homotopic if and only if they are homotopic. Note that Smale-Hirsch theory gives much more general results about homotopy classes of embedded manifolds. We refer the reader for example to Miller's treatment in this specific case [Mil19, Theorem 3.3].

Remark 3.0.2. Two surfaces are regularly homotopic if and only if they differ by a sequence of finger moves and Whitney moves. Moreover, a generic regular homotopy is a sequence of finger moves and Whitney moves. See [FQ90] for a detailed treatment.

3.1 Plan of the chapter.

The proof of Theorem G will be set out as follows.

- **Section 3.2:** We will define surface tubing diagrams, which are similar to those defined by Gabai [Gab17, Definition 5.5]. These contain the data to replace an immersed surface with an embedded one by removing discs around pairs of intersection points and adding in tubes which run along the surface and join up the resulting boundary circles.
- **Sections 3.3 and 3.4:** We prove the *tube swap lemma*, Lemma 3.3.1, and the *tube move lemma*, Lemma 3.4.1. These prove that we can change the way in which the surface is tubed (i.e. change the surface tubing diagram) by performing a single stabilisation followed by a single destabilisation.
- **Section 3.5:** We associate to a regular homotopy a sequence of stabilisations and destabilisations by *shadowing*, again using the terminology of Gabai [Gab17], the homotopy by a tubed surface as follows.
 - At each stage in the regular homotopy where a finger move is performed, we remove the pair of double points by performing a stabilisation which adds a tube running along the surface.

- At each stage in the regular homotopy where a Whitney move is performed, we wish to perform a destabilisation which removes a tube. However, this may not always immediately be possible for several reasons. Firstly, pairs of double points removed by the Whitney move may not have been tubed to each other, but rather to other double points. Secondly, the tubes may not run over one of the arcs of the Whitney circle. Thirdly, the pairing may be ‘crossed’ in the terminology of Gabai [Gab17, §3.3]. We use the results proved in Sections 3.3 and 3.4 to change the way that the surface is tubed to eliminate the first two difficulties. We reduce the third difficulty to the fact that two surfaces $K, K' \subset B^4$ which are slice surfaces for the unknot, both have the property that they destabilise to the standard disc bounded by the unknot.
- **Section 3.6:** We prove that the pair of slice surfaces K and K' obtained in Section 4 both destabilise to the standard disc bounded by the unknot.

3.2 Surface tubing diagrams

We now define surface tubing diagrams which describe how to turn an immersed surface of genus g with $2n$ double points into an embedded surface of genus $g + n$.

Definition 3.2.1. Given an arc Γ in the plane $\mathbb{R}^2 \times 0 \subset \mathbb{R}^2 \times \mathbb{R}^2$, the *linking annulus* of Γ is the annulus $\Gamma \times S^1 \subset \mathbb{R}^2 \times \mathbb{R}^2$.

Definition 3.2.2. Suppose A , B , and C are immersed surfaces in X , which intersect each other and themselves transversely only in double points. Suppose that $a \in A \cap C$ and $b \in B \cap C$. Suppose that Γ is an embedded arc in C (whose interior is disjoint from A and B , and from other double points of C). Then *tubing* along Γ is the result of removing a disc around a from A , a disc around b from B , and adding the linking annulus of Γ (parametrising $\nu(\Gamma) \cong B^4 \subset X$ as $\mathbb{R}^2 \times \mathbb{R}^2$ so that $C \cap \nu(\Gamma)$ is the plane $\mathbb{R}^2 \times 0$), then smoothing corners; see Figure 3.1.

Remark 3.2.3. Usually when we tube, A , B , and C will, in fact, be subsurfaces of the same connected surface in X , obtained by taking the intersection of the surface with some ball $B^4 \subset X$. Note that the resulting surface, in this case, is oriented if and only if the self-intersections a and b have opposite sign.




Figure 3.1: A neighbourhood $\nu(\Gamma)$ in which the tubing operation takes place. The time direction is depicted upwards and we see the sheet of the surface containing Γ continues into the past and future. Tubing along Γ removes two intersections by removing discs from A and B , and adding the linking annulus of the arc Γ .

Remark 3.2.4. By convention, the middle time picture in movies such as Figure 3.1 will be referred to as the $t = 0$ slice or *the present*.

3.2.1 Tubed surfaces

We shall describe how to replace self transverse immersed surfaces of genus g in X , with $2n$ double points, with embedded surfaces of genus $g + n$ by pairing up double points with opposite sign, and tubing along an arc between them. We make the following definition, analogous to that made by Gabai [Gab17, Definition 5.5].




Figure 3.2: A surface tubing diagram. We depict the preimage S of the immersion, and the two preimages of each double point, for example, x_1^+ and y_1^+ are the preimages of the double point $z_1^+ \in X$. We also show the arcs γ_1 and γ_2 , along which we tube to construct the associated tubed surface \tilde{S} .

Definition 3.2.5. Let X be a smooth oriented 4-manifold. A *surface tubing diagram* \mathcal{S} consists of the following data:

- (i) A compact, oriented, connected surface S possibly with boundary.
- (ii) A generic (only isolated double points), self transverse immersion $f: S \rightarrow X$ with $\partial X \cap f(S) = f(\partial S)$ and with the same number of positive self-intersections

as negative. We take care not to confuse the abstract surface S with the immersed surface which is its image, $f(S)$, denoted S^{im} .

- (iii) A partition of the set of double point images, $\{z \in X : |f^{-1}(z)| = 2\}$ into pairs $\{z_i^+, z_i^-\} \subset X$, $i = 1 \dots n$ with the sign of the intersection z_i^\pm being ± 1 . We denote the preimage of z_i^+ by the pair $(x_i^+, y_i^+) \in S \times S$, which comes with a choice of ordering. We denote the preimage of z_i^- by the pair (x_i^-, y_i^-) . We refer to any pre-images of double points as *marked points* of the diagram.
- (iv) A set of disjoint embedded arcs $\gamma_1, \dots, \gamma_n : [0, 1] \rightarrow S$ with endpoints $\gamma_i(0) = x_i^+$ and $\gamma_i(1) = x_i^-$ and which are disjoint from $\{x_j^\pm, y_j^\pm\}_{j \neq i}$ and from y_i^\pm (note that their images $\Gamma_i = f(\gamma_i([0, 1]))$ are also disjoint embedded arcs in S^{im}); see Figure 3.2.

Remark 3.2.6. In the topological case we make the same definition for $f: S \rightarrow X$ which is an immersion obtained from some locally flat embedding $g: S \rightarrow X$ by applying some sequence of finger moves, Whitney moves, and ambient isotopy.

Definition 3.2.7. Given a surface tubing diagram we construct the *associated tubed surface* by tubing S^{im} to itself along each arc Γ_i using tubes in a small neighbourhood of each Γ_i , as in Definition 3.2.2 with $\Gamma = \Gamma_i$, $A = f(\nu(y_i^+))$, $B = f(\nu(y_i^-))$, and $C = f(\nu(\gamma_i))$; again see Figure 3.1. Since z_i^+ and z_i^- have opposite signs the result is an oriented embedded surface which we call \tilde{S} .

Remark 3.2.8.

- (1) Ambient isotopy of the immersed surface S^{im} (which gives rise to an isotopy of the immersion f) describes an ambient isotopy of the arcs Γ_i , and so by extension an isotopy of the associated tubed surface \tilde{S} , where we make sure to keep the tubes close to the surface. At the end of the isotopy, we still have a tubed surface and we still have a surface tubing diagram with the same arcs (but with immersion data $H(-, 1) \circ f$, where $H: S \times [0, 1] \rightarrow X$ is the ambient isotopy).
- (2) An isotopy of the set of arcs γ_i in S (which keeps the arcs disjoint from the preimages of self-intersection points throughout the isotopy) gives rise to an isotopy of the tubes and hence of the associated tubed surface.

3.3 The Tube Swap Lemma

We prove that if we change our surface tubing diagram as depicted in Figure 3.3, then there is a single stabilisation followed by a single destabilisation taking one associated surface to the other.

The proof is mainly by pictures. We shall draw surfaces in a 4-ball as movies with a time direction, drawing slices of 3-dimensional space. For the sake of readability, we draw our pictures as piecewise-smooth and so they have corners, however, they should each be understood to describe a smooth surface. The corners arise in two ways, firstly from stabilisations and destabilisations, and secondly when part of the surface ‘jumps’ into the time direction. The reader should mentally smooth these pictures. For stabilisations and destabilisations, the smoothing is as in Figure 1.1. The corners arising from jumps into the time direction are locally modelled as a product of two arcs which are properly embedded in two discs, one arc having a corner. Smoothing the corner of the arc gives a smoothing of the surface.

Lemma 3.3.1 (Tube Swap Lemma). *Given a surface tubing diagram \mathcal{S} , let β be any arc in S from y_i^+ to y_i^- which is disjoint from all marked points and curves in the surface diagram (including γ_i). Form \mathcal{S}' by removing the arc γ_i and replacing it with β (and changing the order of (x_i^+, y_i^+) to (y_i^+, x_i^+) and (x_i^-, y_i^-) to (y_i^-, x_i^-)); see Figure 3.3.*

Then the associated tubed surface \tilde{S} can be transformed into \tilde{S}' by performing a single stabilisation, followed by a destabilisation (and ambient isotopy).




Figure 3.3: The diagram for a tube swap: we swap the arc γ_i for any arc β from y_i^+ to y_i^- , disjoint from all marked arcs and points. The tube swap lemma says there is a stabilisation and destabilisation, taking the associated tubed surface of the left diagram, to the associated tubed surface of the right.

Proof. We direct the reader to Figure 3.4. We first consider a small tubular neighbourhood $\nu(\Gamma_i \cup f(\beta))$ in X , which is diffeomorphic to $S^1 \times B^3$. We pick a diffeomorphism of $\nu(\Gamma_i)$ to B^4 , parametrising so that at $t = 0$ we see $f(\nu(\gamma_i))$ (the sheet of the immersed

surface containing Γ_i) and the ends of the arc $f(\beta)$). The rest of the sheet containing the ends of $f(\beta)$ extends into the past and future. We extend our parametrisation to one of $S^1 \times B^3$ so that at each t we see a copy of $S^1 \times B^2$, and so $f(\beta)$ is in the $t = 0$ frame and the sheet containing it extends into the past and future; see Figure 3.4.




Figure 3.4: The surface S^{im} in the neighbourhood $\nu(\Gamma \cup f(\beta)) \cong S^1 \times B^3$. The middle picture is $t = 0$. In each time frame shown we see a copy of $S^1 \times D^2$.

The associated tubed surface \tilde{S} in this parametrisation is depicted in Figure 3.5.A. We perform a stabilisation between the top of the tube and the subset $f(\nu(\gamma_i))$ of the surface, to obtain Figure 3.5.B (provided the stabilisation pictured respects orientations; we deal below with the case in which it does not).

To obtain Figure 3.6.C from Figure 3.5.B we perform an isotopy supported in the $t = 0$ frame by ‘inflating’ the tube from the stabilisation. To obtain Figure 3.6.D we flatten the resulting bulge, which slightly deforms the surface in other time frames.

Next, we push the sides of the tube into the $t = 0$ frame to obtain Figure 3.7.E. Then we perform an isotopy, flattening the resulting dip to obtain Figure 3.7.F. To obtain Figure 3.7.G from Figure 3.7.F, we thicken the red arc $f(\beta)$ by pushing some of the surface from the future into the present. We then perform an isotopy to create a tube with a band removed, depicted in Figure 3.7.H.

To obtain Figure 3.8.I from Figure 3.7.H we perform an isotopy. In the same picture, we depict an embedded disc intersecting the surface on its boundary in Figure 3.8.J. We then remove an annulus given by a neighbourhood of the boundary of this disc, and replace it with two parallel copies of the disc in the $t = 0$ frame, thus performing a destabilisation, to obtain Figure 3.8.K. We then push part of the band containing the arc $f(\beta)$ into the future and perform a small further isotopy to obtain Figure 3.8.L. Note that Figure 3.8.L is precisely \tilde{S} .




Figure 3.5: In A we depict the associated tubed surface \tilde{S} in a neighbourhood of X given by $\nu(\Gamma \cup f(\beta)) \cong S^1 \times B^3$. To obtain B from A we perform a stabilisation from the tube to the sheet of the surface $f(\nu(\gamma_i))$.




Figure 3.6: To obtain C from Figure 3.5.B we perform an isotopy which is supported in the $t = 0$ frame ('inflating' the tube from the stabilisation). To obtain D from C we then flatten the picture out, which slightly deforms the surface in other time frames.




Figure 3.7: Here the time direction is to the right, and we depict 4 stages of the isotopy. To obtain E from Figure 3.6.D we push the sides of the tube into the $t = 0$ frame. To obtain F we perform an isotopy to E which flattens the surface. To obtain G from E we then ‘thicken’ the red arc by pushing some of the surface from the future into the $t = 0$ frame. To obtain H from G we then perform an isotopy of this band, to form a tube with a band removed.




Figure 3.8: Here the time direction is to the right, and we depict 4 stages of the isotopy. To obtain I from Figure 3.7.H we perform a small isotopy of the surface. In J we depict a disc intersecting the surface on the boundary of the disc. To obtain K from J we perform a destabilisation by removing a neighbourhood of the boundary of the disc and adding two parallel copies of the disc. To obtain L from K we perform an isotopy (pushing some of the middle band into the future) to obtain the associated tubed surface \tilde{S}' .




Figure 3.9: A stabilisation and isotopy respecting the alternative possible orientations. We only depict the $t = 0$ slice, the past and future frames are as in the previous case, until the final picture where the past and future pictures are swapped. In A we see the associated surface, to which we already performed a stabilisation respecting the alternative orientations. To obtain B from A we perform an isotopy ‘inflating’ the tube from the stabilisation. To obtain B from C we flatten the surface (which deforms the past and future images slightly). To obtain D from C we push the sides of the tube into the present. To obtain E from D we perform an isotopy to flatten the surface. To obtain F from E we perform an isotopy. F is identical to Figure 3.7.F, though we note the past and future images are swapped. From this point we proceed as in the previous case, with the past and future images swapped.

In the case that the above stabilisation was not compatible with orientations, we instead perform a stabilisation on the underside of the surface to obtain Figure 3.9.A. We then perform the sequence of isotopies in Figure 3.9, which are analogous to those in the previous case. Note that in Figure 3.9 only the $t = 0$ slice is depicted, the past and future pictures are as in the previous case, until the final picture where the past and future images are swapped (i.e. the past images become the future images and *vice versa*). Figure 3.9.F is identical to Figure 3.7.F, except that the past and future images are swapped. We now proceed as above (with the past and future images swapped). Again we obtain \widetilde{S}' , completing the proof of the tube swap lemma. \square

3.4 The Tube Move Lemma

We prove that if we take a single arc γ_i in our surface tubing diagram, remove it and replace it with another arc, as in Figure 3.10, there is a single stabilisation and single destabilisation taking one associated surface to the other.




Figure 3.10: A diagram for a tube move. We replace γ_i with α , which may intersect γ_i , but not other marked arcs or points.

Lemma 3.4.1 (Tube Move Lemma). *Given a surface tubing diagram \mathcal{S} , let α be an arc in \mathcal{S} from x_i^+ to x_i^- which is disjoint from the curves $\{\gamma_k\}_{k \neq i}$ (note that α may intersect γ_i), and is also disjoint from all marked points on the surface other than x_i^+ and x_i^- . Form \mathcal{S}' by removing the arc γ_i and replacing it with α .*

Then the associated tubed surface \widetilde{S} can be transformed into \widetilde{S}' by performing a single stabilisation, followed by a destabilisation (and ambient isotopy).

Proof. We direct the reader to Figure 3.11. In Figure 3.11.A, we see the associated tubed surface \widetilde{S} in a neighbourhood of a point $p \in \Gamma_i$ which intersects the linking annulus of Γ_i on a smaller annulus. First, we perform a stabilisation between the tube and the surface




Figure 3.11: A stabilisation and sequence of isotopies allowing us to cut open tubes in order to move their ends about freely. In A we depict a neighbourhood $\nu(p) \subset X$ for some $p \in \Gamma_i$. Note that $\nu(p) \cong B^4$. The time direction is depicted upwards and in each time frame, we see a copy of B^3 . To obtain B from A we perform a stabilisation. To obtain C from B we perform an isotopy in the $t = 0$ frame, which also slightly deforms the surface in other time frames. To obtain D from C we push part of the sides of the tube into the $t = 0$ frame. To obtain E from D we perform an isotopy flattening the surface. Finally, in F we see two disjoint tubes which join the surface at their ends as pictured. These ends may now be pushed about the surface. When we later rejoin the tubes we read the pictures in reverse order.

to obtain Figure 3.11.B (note that if this is not compatible with orientations, we instead perform the stabilisation on the underside and proceed in the same way, taking the mirror image of each figure). We then perform an isotopy to obtain Figure 3.11.C, then push part of the sides of the tube into the present to obtain Figure 3.11.D. We then perform a further isotopy to obtain Figure 3.11.E. Here we see two disconnected tubes, whose ends join the surface as depicted in Figure 3.11.F.

We can now move these ends around freely on the surface, provided during the isotopy they are disjoint from the other tubes and double points. We depict how we drag the ends diagrammatically in Figure 3.12. First, we retract the two tubes along Γ_i dragging the ends towards z_i^+ and z_i^- . We then perform an isotopy that pushes these ends along α so that the two tubes now run along α and the ends of the tubes lie in a neighbourhood of some point $q \in f(\alpha)$. In this neighbourhood, we see the final picture of Figure 3.11. We then rejoin the tubes, by performing an isotopy and destabilisation which can be seen by reading the pictures in Figure 3.11 in reverse order. After rejoining the tubes they form the linking annulus of α . The resulting surface is \widetilde{S} as required. This completes the proof of the tube move lemma. \square




Figure 3.12: A schematic of the proof of the tube move lemma. The second image depicts the cut-open tube. The crosses depict where the ends of the cut-open tube join the surface. We then show how to move these ends in the next two equivalences. To obtain to the final picture we rejoin the cut open ends.

Remark 3.4.2. Note that we cannot prove Lemma 3.4.1 by using Lemma 3.3.1 twice, due to the condition in Lemma 3.3.1 that the new arc is disjoint from the old one. To do so we would need to find an intermediate arc β from y_i^+ to y_i^- disjoint from both γ_i and α , but if γ_i and α intersect, then $S \setminus \{\gamma_i, \alpha\}$ may be disconnected so this may not be possible.

3.5 Proof of Theorem G

With these tools in place, we proceed with the proof of Theorem G, namely that $d_{\text{st}}(\Sigma, \Sigma') \leq d_{\text{sing}}(\Sigma, \Sigma') + 1$. We do so by shadowing a regular homotopy by a sequence of embedded surfaces differing by stabilisation, destabilisation, and ambient isotopy.

Proof of Theorem G. Recall that Σ and Σ' are embedded surfaces in X of genus g . In the case the surfaces are not regularly homotopic $d_{\text{sing}}(\Sigma, \Sigma') = \infty$ and we are done. Hence assume that Σ and Σ' are regularly homotopic and suppose $d_{\text{sing}}(\Sigma, \Sigma') = n$. Given a distance minimising regular homotopy from Σ to Σ' , let P_1, \dots, P_k be the sequence of immersed surfaces describing this homotopy, each differing from the previous by either a finger move, a Whitney move, or an ambient isotopy.

We shall describe a sequence of surface tubing diagrams with associated tubed surfaces that differ by stabilisations, destabilisations, and ambient isotopy, such that the genus of any intermediate surface never exceeds $g + n + 1$.

Fix the abstract surface S , and immersions $f_i: S \rightarrow X$ with image P_i for each i . The immersion $f_1: S \rightarrow X$ gives a surface diagram \mathcal{S}_1 (the empty diagram in S) with associated tubed surface $\tilde{\mathcal{S}}_1 = S_1^{\text{im}} = P_1$.

We now suppose for induction that we have a surface tubing diagram \mathcal{S}_i with immersion data $f_i: S \rightarrow X$, and associated tubed surface $\tilde{\mathcal{S}}_i$. We will construct a surface tubing diagram \mathcal{S}_{i+1} with immersion data $f_{i+1}: S \rightarrow X$, such that $\tilde{\mathcal{S}}_{i+1}$ differs from $\tilde{\mathcal{S}}_i$ by a series of stabilisations, destabilisations, and ambient isotopy, such that the genus of any intermediate surface does not exceed $g + n + 1$. There are three cases; P_{i+1} is obtained from P_i by ambient isotopy, a finger move, or a Whitney move.

Ambient Isotopy: If P_{i+1} just differs from P_i by ambient isotopy, by Remark 3.2.8(1) we have a new surface tubing diagram \mathcal{S}_{i+1} with immersion data $f_{i+1}: S \rightarrow X$. Furthermore, $\tilde{\mathcal{S}}_i$ is ambiently isotopic to $\tilde{\mathcal{S}}_{i+1}$.

Finger Move: If P_{i+1} differs from P_i by a finger move, let α and β be the Whitney arcs of the Whitney disc that undoes the finger move. Note these arcs are in S , the abstract surface. Before we perform the finger move, we perform an isotopy of the arcs, to push any arcs off α and β to form $\mathcal{S}'_i = (f_i, \{\gamma'_j\})$. We then add the new double points and β to the diagram to obtain the diagram $\mathcal{S}_{i+1} = (f_{i+1}, \{\gamma'_j\} \cup \{\beta\})$ which we note uses the new

map f_{i+1} , which differs from f_i by the finger move; see Figure 3.13. The associated tubed surfaces \widetilde{S}_i and \widetilde{S}'_i are ambiently isotopic by Remark 3.2.8(2), and \widetilde{S}'_i and \widetilde{S}_{i+1} differ by isotopy and a stabilisation, as in Figure 3.14.




Figure 3.13: The sequence of surface tubing diagrams corresponding to performing a finger move.




Figure 3.14: Above we see the immersed surface before and after a finger move. Below we perform a sequence of isotopies and stabilisations taking \widetilde{S}'_i to \widetilde{S}_{i+1} . First, we isotope one sheet to create a Whitney bubble and push the other sheet into this bubble. We then stabilise the bubble to create the linking annulus of Γ_i .

Whitney Move: The Whitney moves present the main difficulty. If P_{i+1} differs from P_i by a Whitney move, there are several cases to consider.

The endpoints of the Whitney arcs form the set $\{x_i^+, y_i^+, x_j^-, y_j^-\}$ for some i and j . In the case that $i = j$ then either $\partial\alpha = \{x_i^+, x_i^-\}$ and $\partial\beta = \{y_i^+, y_i^-\}$, or $\partial\alpha = \{x_i^+, y_i^-\}$ and $\partial\beta = \{y_i^+, x_i^-\}$ (up to relabeling of α and β). We call the former situation Case 1 and the latter situation Case 2. Case 2 is the crossed Whitney disc .

Otherwise, if $i \neq j$ then either $\partial\alpha = \{x_i^+, x_j^-\}$ and $\partial\beta = \{y_i^+, y_j^-\}$, or $\partial\alpha = \{x_i^+, y_j^-\}$ and $\partial\beta = \{y_j^+, x_i^-\}$ (again up to relabeling of α and β). We call the former situation Case

3 and the latter Case 4.

Note that α and β may intersect other arcs, including γ_i and γ_j . In all cases, we perform a small isotopy of all the arcs so that they intersect α and β transversely.

Case 1: *The Whitney arcs are α between x_i^+ and x_i^- , and β between y_i^+ and y_i^- for some i . We wish to run in reverse what we did in the case of a finger move, however, we need γ_i to run over either α or β , and we also need α and β to be disjoint from other arcs, so that the surface tubing diagram looks like the diagram at the end of a finger move, as in Figure 3.13.*




Figure 3.15: The sequence of tubing diagrams which correspond to performing a Case 1 Whitney move. First, we move all arcs off β , then swap γ_i to β , then move all arcs off α . The resulting diagram corresponds to that at the end of a finger move. We can now destabilise the associated surface to obtain the associated surface for the final diagram (which uses the new immersion data), by running Figure 3.14 backwards.

We arrange this by first using the tube move lemma to move any arcs off β . To do so we remove intersection points with β one by one. We consider the arc γ_r which has the intersection with β closest to y_i^+ (in the sense of distance along β), at $p \in \beta$. We form the arc γ'_r , by removing an arc neighbourhood of p from γ_r , and replacing it with an arc that runs along the boundary of a small neighbourhood of $\beta \subset S$; see Figure 3.15. Provided the neighbourhood is taken to be sufficiently small, γ'_r is disjoint from $\{\gamma_p\}_{p \neq r}$ and from marked points other than x_r^\pm , and so the corresponding associated surfaces differ by a stabilisation and destabilisation by the tube move lemma. We do this for every intersection point of arcs with β , until all arcs are disjoint from β . Note that we may move γ_i during this process.

Next, we use the tube swap lemma to swap γ_i to β . We then use the tube move lemma to

move any arcs off α as we did for β ; see the third to fourth picture of Figure 3.15. Finally, we perform the destabilisation and isotopy coming from reading Figure 3.14 in reverse. This has the effect of removing the points x_i^\pm and y_i^\pm , and the arc γ_i (which now runs along β) from the diagram.




Figure 3.16: The sequence of diagrams corresponding to a Case 2 Whitney move. First, we remove all arcs running over β and α using the tube move lemma, then claim we may remove the resulting tube and intersection points using a stabilisation and two destabilisations.

Case 2: The Whitney arcs are α between x_i^+ and y_i^- , and β between y_i^+ and x_i^- for some i . This is a crossed Whitney disc. As in Case 1, we first remove any arcs running over α and β using the tube move lemma, which may move γ_i ; see Figure 3.16. Let S'_i be the resulting surface diagram, and S_{i+1} be the diagram with the same points and arcs but with the points x_i^\pm , y_i^\pm , the arc γ_i deleted, and immersion data f_{i+1} ; see Figure 3.16.

We show \widetilde{S}'_i and \widetilde{S}_{i+1} are related by a stabilisation and two destabilisations (and isotopy). We depict the immersed surface S'^{im}_i in a neighbourhood of the Whitney disc in Figure 3.17. The associated tubed surface in this neighbourhood is shown in Figure 3.18. To obtain the second picture of Figure 3.18 we perform an isotopy which pulls apart the surface (taking the sheets of the surface to where they need to be after the Whitney move), at the expense of creating a *double tube* again using the terminology of Gabai [Gab17], a Hopf link $\times [0, 1]$ running through X between two disc neighbourhoods of the surface. The Hopf link $\times [0, 1]$ has four boundary components, two of which are joined to the surface (the top of the green tube and bottom of the pink), while the other two join to the linking annulus of Γ_i (the bottom of the green and top of the pink).

We now perform a stabilisation inside the Hopf link $\times [0, 1]$, which can be seen in the movie picture as attaching two bands (provided this is compatible with orientation, if not see below); see Figure 3.19. The middle picture of this movie is then an unknot, along which




Figure 3.17: The immersed surface in a neighbourhood of a crossed Whitney disc. The arc Γ_i joins the two sheets, running over the surface outside of this 4-ball.




Figure 3.18: The associated tubed surface \tilde{S}_i . We pull the two sheets of the associated tubed surface apart, at the expense of creating a double tube. Note the pink and green tubes join up outside of this picture, and form the linking annulus for Γ_i . This is the same operation as that Gabai describes in [Gab17, Figure 5.9].




Figure 3.19: The first image depicts part of the double tube. To obtain the second we perform a stabilisation. We then perform an isotopy to obtain the third, in which we see an unknot $\times [0, 1]$ in the middle time frame. To obtain the fourth we perform a destabilisation along this unknot.




Figure 3.20: The surface in a neighbourhood of Γ_i . On the left we depict part of the double tube; we show how the double tube joins up using the dotted arcs (though this does not happen in this 4-ball). On the right, we depict part of the slice surface K ; we do not draw K inside the two 4-balls indicated. Inside these two 4-balls the surface is given by standard annulus bounded by the Hopf link, whose interior has been pushed into the 4-ball. To obtain the surface on the right from that on the left, we cut open the double tube as in Figure 3.19 and retract the two capped off ends so that they lie near the surface, inside the neighbourhood of Γ_i .

we perform a destabilisation to cut the double tube; again see Figure 3.19.

The resulting ‘end’ of the double tube can be made by taking an open double tube, adding a band between the two tubes, then capping off the resulting boundary circle with a disc. It can also be thought of as the standard annulus whose boundary is the Hopf link whose interior is pushed into the 4-ball. We now suck back these ends to be close to the surface. The resulting surface differs from \tilde{S}_{i+1} only in a 4-ball neighbourhood of the arc Γ_i . The resulting capped off double tube in this 4-ball neighbourhood is pictured on the right in Figure 3.20. In this neighbourhood, we see a genus 1 slice surface for the unknot. We call this slice surface K .

If the above stabilisation was not compatible with the orientation on the double tube, we instead perform a different stabilisation and destabilisation which are compatible with the alternative possible orientations; see Figure 3.21. Again the result is a cut open double tube, and we again suck the ends back to lie in a 4-ball neighbourhood of Γ_i . We call the resulting surface in this neighbourhood $K' \subset B^4$ and note the boundary of K' is the unknot in S^3 as before.




Figure 3.21: Cutting the double tube using a stabilisation and a destabilisation which are compatible with the other possible orientation of the double tube. We perform an isotopy of the surface to obtain the second picture which shows the result differs by just a twist in the bands.

We now wish to construct a sequence of stabilisations and destabilisations taking K and K' to the trivial disc. Such a sequence would remove the mess of tubes, and take our surface to \tilde{S}_{i+1} as required. We exhibit such a sequence in Lemma 3.6.2. Note that results in [BS15] imply that some sequence of stabilisations and destabilisations exists but we show that in fact, one destabilisation is sufficient, and so the genus of any intermediate surface does not exceed $g + n + 1$.

Case 3: *The Whitney arcs consist of an arc α between x_i^+ and x_j^- , and an arc β between*

y_i^+ and y_j^- for $i \neq j$. This case presents the most diagrammatic difficulty. We wish to remove the intersections of arcs with α and β , however this is made difficult by the fact that γ_i and γ_j may intersect α and β in a complicated way; see Figure 3.22.




Figure 3.22: The sequence of diagrams corresponding to a Case 3 Whitney move. We first pick arcs a_i and a_j disjoint from α and β . To obtain the second diagram we use the tube move lemma to remove intersections of any arcs with the arc $a_i \cup \beta \cup a_j$. To obtain the third diagram we use the tube swap lemma twice. We then remove intersections with α using the tube move lemma. Finally, we remove the points removed the Whitney move and join the tubes; the corresponding destabilisation is pictured in Figure 3.23. Note that the final diagram uses the new immersion data.

To overcome this, we first we pick an arc a_i from y_i^+ to y_i^- which is disjoint from $\alpha \cup \beta$ and all other marked points on the surface; clearly, we may do so since the complement of $\alpha \cup \beta$ and all points is connected. We then similarly pick an arc a_j from y_j^+ to y_j^- which is disjoint from $\alpha \cup \beta \cup a_i$ and all marked points other than y_j^+ and y_j^- ; again the complement of these arcs and points is connected since a_i does not intersect α or β . We remove the intersections of all arcs with $a_i \cup \beta \cup a_j$, which is one long embedded arc, one by one using the tube move lemma as in previous cases; see Figure 3.22. Note that this may move γ_i and γ_j .

We now use the tube swap lemma to swap γ_i to a_i and γ_j to a_j . We then move any arcs

off α using the tube move lemma.

Finally we remove the intersection points $y_i^+, x_i^+, y_j^-, x_j^-$, join the arcs γ_i and γ_j using β , and change the index of the points labelled j to i ; see Figure 3.22. The corresponding isotopy and destabilisation of the associated surface is given by Figure 3.23.




Figure 3.23: Above we see the Whitney move. Below we see the corresponding isotopy and destabilisation of the associated surfaces. This operation is the operation described by Gabai in [Gab17, Figure 5.8] with an additional destabilisation to remove the ‘single tube’ pictured in [Gab17, Figure 5.8].

Case 4: *The Whitney arcs consist of an arc α between x_i^+ and y_j^- , and an arc β between y_i^+ and x_j^- for $i \neq j$. In this case, we use the tube swap lemma to swap γ_i to any arc between y_i^+ and y_j^- disjoint from other tube arcs, but which may intersect α and β ; some such arc exists since the complement of $\cup_r \gamma_r \cup_k y_k^\pm$ is a punctured surface, so is connected. We are then in Case 3 and proceed as before.*

At stage P_k , we have a surface diagram \mathcal{S}_k with immersion data $f_k \rightarrow X$. Since f_k is an embedding it has no intersection points so must be the empty diagram, hence the associated tubed surface \tilde{S}_k is $f_k(S) = P_k$ as required. This completes the proof of Theorem G modulo Lemma 3.6.2. □

Examining the proof, we in fact prove a stronger, if more technical, fact.

Proposition 3.5.1. *Let Σ and Σ' be immersed surfaces with $|\text{sing}(\Sigma)| = |\text{sing}(\Sigma')|$, both with the same number of positive and negative double points. Suppose they are regularly homotopic through surfaces with at most $2n$ double points. Then given a surface tubing diagram \mathcal{S} for Σ , there exists a surface tubing diagram \mathcal{S}' for Σ' such that*

$$d_{\text{st}}(\tilde{S}, \tilde{S}') \leq n - \frac{1}{2}|\text{sing}(\Sigma)| + 1.$$

Note that taking both Σ and Σ' to be embeddings yields Theorem G, since the only surface tubing diagrams for embeddings are the empty diagram.

Remark 3.5.2.

- (1) Schwartz [Sch18] constructs pairs of embedded spheres in a 4-manifold X with 2-torsion in $\pi_1(X)$, which are regularly homotopic, and are such that any regular homotopy between them must contain a crossed Whitney move.
- (2) For each finger move we made a choice of Whitney arc to tube along; we tubed along β , but we could equally have tubed along α . In the absence of a crossed disc we may make this choice so that only Case 1 and Case 3 Whitney moves occur. Indeed considering a homotopy as a map $H: S \times [0, 1] \rightarrow X \times [0, 1]$, the set of double points is a union of circles. Each double point circle C has two disjoint circles as its preimage, $H^{-1}(C) = C_x \cup C_y$. Labelling the double point preimages so that $x_i^\pm \in C_x$ and $y_i^\pm \in C_y$ gives such a choice of tubing.

3.6 Destabilising the slice surfaces K and K'

We complete the proof of Theorem G by showing that both K and K' become the standard disc bounded by an unknot in S^3 after a single destabilisation.

3.6.1 Banded link presentations of knotted surfaces

First, we review the calculus of banded link presentations for slice discs and 2-knots set out by Jablonowski [Jab16].

Definition 3.6.1. A *banded link presentation* of a smooth embedded surface $\Sigma \subset B^4$ bounded by a link L in S^3 , consists of the link $L' = L \cup U_n$, where U_n denotes the unlink with n components, along with a number of bands, i.e. embedded copies $[0, 1] \times [0, 1]$ disjoint from each other and L' , except at the ends of the bands $\{0, 1\} \times [0, 1]$ which lie in L' . Performing a *band move* is the operation of removing the ends of the bands $\{0, 1\} \times [0, 1]$ from L' and adding in the sides of the bands $[0, 1] \times \{0, 1\}$. We require that the resulting link after performing all the band moves to L' is the unlink.




Figure 3.24: Moves on banded link diagrams giving isotopic surfaces. The top left equivalence is a band slide, the top right equivalence a band swim. The bottom left equivalence is the cancellation of a maximum and a saddle, the bottom right equivalence is the cancellation of a minimum with a saddle.

A banded link presentation describes a slice surface for L via a movie which can be seen by considering B^4 as $B^3 \times [0, 1]$. The first slide in this movie is L . In the next, we add the unknotted components of U_n , these correspond to minima of the surface with respect to the projection onto $[0, 1]$. In the next slide we perform a band move on each band, which corresponds to adding saddles to the surface. After performing these band moves we obtain the unlink $U_{n'}$ for some n' . Each component of this unlink is then capped off with a disc, corresponding to maxima of the surface.





Figure 3.25: Stabilisation and destabilisation in banded link presentations.

There are several moves on banded link diagrams one may perform that give isotopic surfaces. These are isotopy of the diagram, band slides, band swims, cancelling a maximum with a saddle, and cancelling a minimum with a saddle; see Figure 3.24.

Stabilisation or destabilisation in banded link diagrams corresponds to adding or removing respectively two bands, as in Figure 3.25. Note that it does not matter where the loose end of the band goes.

3.6.2 Proof that K and K' destabilise to the standard disc

Lemma 3.6.2. *The surfaces K and K' , described in Case 2 in the proof of Theorem G, become the standard disc bounded by the unknot in S^3 after a single destabilisation.*



Proof. We first consider Figure 3.20. We recall that the ‘capped-off’ ends are made by attaching a band between the tubes and capping off by a disc as in Figures 3.19 and 3.21. After an isotopy, the z -direction as depicted in Figure 3.26 gives the standard Morse function for B^4 , which restricts to a Morse function on the surface. With this Morse function the surface has one minimum and two saddles, then two further saddles and two maxima from the caps.




Figure 3.27: A movie for the slice surface K . Note that the z direction depicted in Figure 3.26, is now the time direction of our movie.

This Morse function gives a movie presentation for the surface in the 4-ball; see Figure 3.27. We deform this into a band presentation, depicted in Figure 3.28, which we simplify using band swims, slides, and a destabilisation, to obtain the standard disc.

In the case of K' , recall that we stabilised the outside of one tube with the inside of the other, which has the effect of adding a half twist to the bands; see Figure 3.29. After a band swim and isotopy, we obtain the mirror image of Figure 3.28 and proceed as before (taking the mirror image of each picture), destabilising to obtain the standard disc. \square




Figure 3.28: A band presentation for K , which we destabilise to obtain the standard disc. To obtain the second image from the first we perform an isotopy, untwisting the two bands at the sides. To obtain the third from the second we perform two band slides and one band swim, as indicated by the arrows. To obtain the fourth we perform a band slide as indicated, and an isotopy of bands. To obtain the fifth we cancel a minimum with a saddle. To obtain the sixth we perform the indicated band slide, we then perform another band slide to obtain the seventh. We then destabilise to obtain the eighth image. Finally we cancel a maximum and a saddle to obtain the banded link presentation which is just the unknot, which is a banded link presentation for the standard disc bounded by the unknot.




Figure 3.29: A band presentation for K' , which we see is the mirror image of the band presentation for K .

Bibliography

- [BK20] D. Baraglia and H. Konno, *A gluing formula for families Seiberg-Witten invariants*, *Geom. Topol.* 24, no. 3, 1381–1456, 2020.
- [BFL14] A. Bartels, F. T. Farrell, and W. Lück, *The Farrell-Jones Conjecture for cocompact lattices in virtually connected Lie groups.*, *J. Amer. Math. Soc.*, 27(2):339–388, 2014.
- [BS15] R. I. Baykur and N. Sunukjian, *Knotted surfaces in 4-manifolds and stabilizations*, *J. Topol.* 9, no. 1, pages 215-231, 2015.
- [BG21] R. Budney and D. Gabai, *Knotted 3-balls in S^4* , preprint, arXiv:1912.09029, 2021.
- [Cer70] J. Cerf, *La stratification naturelle des espaces de fonctions différentiables réelles et le théorème de la pseudoisotopie*, *Publ. Math. I.H.E.S.* 39, 1970.
- [DK01] J. Davis and P. Kirk, *Lecture Notes in Algebraic Topology*, No. 35, Graduate Studies in Mathematics. Amer. Math. Soc., 2001.
- [FQ90] M. Freedman, and F. Quinn, *Topology of 4-Manifolds*, Princeton Univ. Press, 1990.
- [Gab17] D. Gabai, *The 4-Dimensional Light Bulb Theorem*, *J. Amer. Math. Soc.*, vol. 33, no. 3, 2020.
- [Hat02] A. Hatcher, *Algebraic Topology*, Cambridge University Press, 2002.
- [HW73] A. Hatcher and J. B. Wagoner, *Pseudo-isotopies of compact manifolds, Part I*, *Astérisque*, tome 6, 1973.
- [Hat73] A. Hatcher, *Pseudo-isotopies of compact manifolds, Part II*, *Astérisque*, tome 6, 1973.
- [Hem76] J. Hempel, *3-Manifolds*, Princeton University Press, *Ann. of Math. Studies*, No. 86, 1976.
- [Hir59] M. W. Hirsch, *Immersions of manifolds*, *Trans. Amer. Math. Soc.* 93, No. 2, pages 242-276, 1959.
- [Igu84] K. Igusa, *What happens to Hatcher and Wagoner’s formula for $\pi_0 C(M)$ when the first Postnikov invariant of M is nontrivial?*, Bak, A. (ed.) *Algebraic K-Theory, Number Theory, Geometry and Analysis, Proceedings of Bielefeld (1982)*. *Series Lecture Notes in Mathematics*, vol. 1046. Springer, 1982.
- [Igu21a] K. Igusa, *Second obstruction to pseudoisotopy I*, preprint, arXiv:2110.09659, 2021.
- [Igu21b] K. Igusa, *Second obstruction to pseudoisotopy in dimension 3*, preprint, arXiv:2112.08293, 2021.
- [Jab16] M. Jablonowski, *On a banded link presentation of knotted surfaces*, *J. Knot Theory Ramifications* 25, No. 3, 2016.

- [Jah] B. Jahren, *On the Wh_1^+ -obstruction and Pseudo-isotopies of Manifolds of dimension 3 and 4*, private correspondence.
- [JZ18] A. Juhász and I. Zemke, *Stabilization distance bounds from link Floer homology*, preprint arXiv:1810.09158v1, 2018.
- [Kre79] M. Kreck, *Isotopy classes of diffeomorphisms of $(k-1)$ -connected almost-parallelizable $2k$ -manifolds*, Algebraic topology, Aarhus 1978 (Proc. Sympos., Univ. Aarhus, Aarhus, 1978), Lecture Notes in Math., vol. 763, Springer, pages 643–663, 1979.
- [KM20] P. Kronheimer and T. Mrowka, *The Dehn twist on a sum of two $K3$ surfaces*, Math. Res. Lett., 27(6):1767–1783, 2020.
- [Kwa87] S. Kwasik, *Low-dimensional concordances, Whitney towers and isotopies*, Math. Proc. of Cambridge Philos. Soc., vol. 102, no. 1, pages 103–119. Springer, 1987.
- [LS69] R. Lashof & J. Shaneson, *Classification of knots in codimension two*, Bull. Amer. Math. Soc., 75, pages 171–175, 1969.
- [Lin20] J. Lin, *Isotopy of the Dehn twist on $K3\#K3$ after a single stabilization*, preprint, arXiv:2003.03925, 2020.
- [LR05] W. Lück and H. Reich, *The Baum-Connes and the Farrell-Jones conjectures in K and L -theory*, Handbook of K-theory, vol. 1, 2, pages 703–842, Springer, 2005.
- [Mil19] M. Miller, *A concordance analogue of the 4-dimensional light bulb theorem*, Int. Math. Res. Not., issue 4, pages 2565–2587, 2021.
- [Qui86] F. Quinn, *Isotopy of 4-manifolds*, J. Differential Geom., vol. 24, No. 3, 1986.
- [Rub98] D. Ruberman., *An obstruction to smooth isotopy in dimension 4*, Math. Res. Lett. 5, no. 6, pages 743–758, 1998.
- [Rub99] D. Ruberman., *A polynomial invariant of diffeomorphisms of 4-manifolds*, In Proceedings of the Kirbyfest (Berkeley, CA, 1998), vol. 2 of Geom. Topol. Monogr., pages 473–488. Geom. Topol. Publ., 1999.
- [Sat69] H. Sato, *Diffeomorphism group of $S^p \times S^q$ and exotic spheres*, Quart. J. Math., Oxford Ser. (2), vol. 20, pages 255–276, 1969.
- [ST19] R. Schneiderman and P. Teichner, *Homotopy versus isotopy: spheres with duals in 4-manifolds*, preprint arXiv:1904.12350, 2019.
- [Sch18] H. R. Schwartz, *Equivalent non-isotopic spheres in 4-manifolds*, J. Topol., vol. 12, pages 1396–1412, 2019.
- [Sin20] O. Singh, *Distances between surfaces in 4-manifolds*, J. Topol., vol. 13, pages 1034–1057, 2020.
- [Sin21] O. Singh, *Pseudo-isotopies and diffeomorphisms of 4-manifolds*, preprint, arXiv:2111.15658, 2021.
- [Sma57] S. Smale, *A classification of immersions of the two-sphere*, Trans. Amer. Math. Soc. 90, No. 2 (1957), 281–290.

-
- [Wat20] T. Watanabe, *Theta-graph and diffeomorphisms of some 4-manifolds*, preprint, arXiv:2005.09545, 2020.
- [Wei05] C. Weibel, *Algebraic K-theory of rings of integers in local and global fields*, Handbook of K-theory, vol. 1, pages 139–190. Springer, 2005.
- [Whi78] G. Whitehead, *Elements of Homotopy Theory*, Graduate Texts in Mathematics No. 61, Springer-Verlag, 1978.