



Geometry & Topology

Volume 28 (2024)

Embedding surfaces in 4-manifolds

DANIEL KASPROWSKI

MARK POWELL

ARUNIMA RAY

PETER TEICHNER

Embedding surfaces in 4–manifolds

DANIEL KASPROWSKI

MARK POWELL

ARUNIMA RAY

PETER TEICHNER

We prove a surface embedding theorem for 4–manifolds with good fundamental group in the presence of dual spheres, with no restriction on the normal bundles. The new obstruction is a Kervaire–Milnor invariant for surfaces and we give a combinatorial formula for its computation. For this we introduce the notion of band characteristic surfaces.

57K40, 57N35

1. Introduction	2399
2. Generic immersions and intersection numbers	2410
3. Secondary embedding obstructions	2429
4. The proof of the surface embedding theorem	2433
5. Band characteristic maps and the combinatorial formula	2437
6. Homotopy versus regular homotopy	2447
7. Proofs of statements from Section 5	2449
8. Proof of Theorems 1.6 and 1.9	2467
9. Examples and applications	2472
References	2479

1 Introduction

We study whether a given map of a surface to a topological 4–manifold is homotopic to an embedding. Here and throughout the article, embeddings and immersions in the topological category are by definition *locally flat*, meaning they are locally modelled on linear inclusions $\mathbb{R}^2 \hookrightarrow \mathbb{R}^4$ or $\mathbb{R}_+^2 \hookrightarrow \mathbb{R}^4$.

Even for maps of 2–spheres, this question has only been completely addressed in a handful of simple manifolds, such as S^4 , $\mathbb{C}\mathbb{P}^2$ [Tristram 1969, page 264] and $S^2 \times S^2$ [Tristram 1969, Theorem 4.5; Kervaire and Milnor 1961, Corollary 1; [Freedman 1982, Corollary 1.1]]. Lee and Wilczyński [1990; 1997] and

Hambleton and Kreck [1993] described the minimal genus of an embedded surface in a fixed homology class, in any given simply connected, closed 4–manifold, assuming that the fundamental group of the complement is abelian. This was recently extended by [Feller et al. 2021] to knot traces. In the relative setting even the simplest case is open: which knots in S^3 bound a (locally flat) embedded disc in D^4 ?

The main available tool for proving positive results is Freedman’s embedding theorem (Theorem 4.3), which shows that maps of discs and spheres to a 4–manifold with *good* fundamental group, with vanishing intersection and self-intersection numbers, and with framed algebraically dual spheres, are regularly homotopic to embeddings [Freedman 1982; Freedman and Quinn 1990, Corollary 5.1B]; see also [Powell et al. 2020; Behrens et al. 2021]. Surgery and the s –cobordism theorem for topological 4–manifolds with good fundamental group are consequences of this theorem [Quinn 1982; Freedman and Quinn 1990]. Our aim, realised in Theorems 1.2 and 1.6 below, is to extend Freedman’s theorem to all compact surfaces with algebraically dual spheres, in any connected 4–manifold with good fundamental group. In Section 1.4, we explain how to apply this to the question from the opening paragraph of whether a given homotopy class contains an embedding. In Section 1.5, we give some applications to knot theory. In particular, we show that every knot bounds an embedded surface of genus one in $M \setminus \mathring{D}^4$ for every simply connected closed 4–manifold M not homeomorphic to S^4 . Recall that, for $M = S^4$, this does not hold because the slice genera of knots can be arbitrary large.

Throughout the paper, we will work in the following setting unless otherwise specified.

Convention 1.1 We assume that M is a connected, topological 4–manifold and that Σ is a nonempty compact surface with connected components $\{\Sigma_i\}_{i=1}^m$. The notation $F = \{f_i\}_{i=1}^m : (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ represents a generic immersion (Definition 2.3) with components $f_i : (\Sigma_i, \partial\Sigma_i) \looparrowright (M, \partial M)$.

By assumption, the map F restricts to an embedding on $\partial\Sigma$ and $F^{-1}(\partial M) = \partial\Sigma$, where $\partial\Sigma$ and ∂M are permitted to be empty. There is no requirement for Σ or M to be orientable, and M could be noncompact. Weakening the hypotheses of Freedman’s theorem to allow for the algebraically dual spheres to be unframed introduces an additional obstruction, the *Kervaire–Milnor invariant* $\text{km}(F) \in \mathbb{Z}/2$ (Definition 1.4), which vanishes in the presence of framed algebraically dual spheres.

Theorem 1.2 (Surface embedding theorem) *Let $F = \{f_i\}_{i=1}^m : (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. Suppose that $\pi_1(M)$ is good and that F has algebraically dual spheres $G = \{g_i\}_{i=1}^m \subseteq \pi_2(M)$. Then the following statements are equivalent:*

- (i) *The intersection numbers $\lambda(f_i, f_j)$ for all $i < j$, the self-intersection numbers $\mu(f_i)$ for all i , and the Kervaire–Milnor invariant $\text{km}(F) \in \mathbb{Z}/2$ all vanish.*
- (ii) *There is an embedding $\bar{F} = \{\bar{f}_i\}_{i=1}^m : (\Sigma, \partial\Sigma) \hookrightarrow (M, \partial M)$, regularly homotopic to F relative to $\partial\Sigma$, with geometrically dual spheres $\bar{G} = \{\bar{g}_i : S^2 \looparrowright M\}_{i=1}^m$ such that $[\bar{g}_i] = [g_i] \in \pi_2(M)$ for all i .*

Equivariant intersection and self-intersection numbers of immersed discs and spheres have a long history (see eg [Wall 1970]). In the theorem above, we consider generalised versions for arbitrary compact surfaces,

lying in quotients of the group ring $\mathbb{Z}[\pi_1(M)]$, which we denote by $\Gamma_{f_i, f_j} \ni \lambda(f_i, f_j)$ for the intersection numbers and $\Gamma_{f_i} \ni \mu(f_i)$ for the self-intersection numbers. We describe these quotients in detail in Sections 2.2 and 2.3. As in the simply connected case, these invariants require based maps (Definition 2.11) but their vanishing as in Theorem 1.2(i) does not depend on the choice of basing. Vanishing of all the $\lambda(f_i, f_j)$ for $i < j$ and all the $\mu(f_i)$ is equivalent to the vanishing of the self-intersection number $\mu(F)$, which is defined as follows.

Definition 1.3 Let $F = \{f_i\}_{i=1}^m : (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. Assume in addition that M and Σ are based and that F is a based map. The *self-intersection number* of this possibly disconnected immersed surface is given by counting all double points of F , as follows:

$$\mu(F) := \sum_{i < j} \lambda(f_i, f_j) + \sum_i \mu(f_i) \in \bigoplus_{i < j} \Gamma_{f_i, f_j} \oplus \bigoplus_i \Gamma_{f_i}.$$

The self-intersection number $\mu(F)$ is a regular homotopy invariant that vanishes if and only if there is a collection of Whitney discs that pair all double points of F (Corollary 2.30), just like for connected surfaces. The Whitney discs may be chosen to form a *convenient* collection, meaning that all Whitney discs have disjointly embedded boundaries, are framed and have interiors transverse to F (Definition 2.31).

Definition 1.4 Let $F : (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. By definition, $\text{km}(F) \in \mathbb{Z}/2$ vanishes if and only if, after finitely many finger moves taking F to some F' , there is a convenient collection of Whitney discs pairing all the double points of F' and whose interiors are disjoint from F' .

We think of $\mu(F)$ as the primary embedding obstruction, and $\text{km}(F)$ as the secondary embedding obstruction. Note that $\text{km}(F) = 0$ implies $\mu(F) = 0$ but $\text{km}(F)$ is always defined even if $\mu(F) \neq 0$. In Section 1.1, we will give a combinatorial description of $\text{km}(F)$ in the case that $\mu(F) = 0$.

The Kervaire–Milnor invariant is named in homage to [Kervaire and Milnor 1961], in which the authors defined an embedding obstruction and used it to give the first proof that the Whitney trick fails in dimension 4. Section 3 gives details on the connection of our Kervaire–Milnor invariant to the original obstruction and other secondary embedding obstructions in the literature.

A group is said to be *good* if it satisfies the π_1 -null disc property [Freedman and Teichner 1995] (see also [Kim et al. 2021]); we shall not repeat the definition. In practice, it suffices to know that virtually solvable groups and groups of subexponential growth are good, and that the class of good groups is closed under taking subgroups, quotients, extensions and colimits [Freedman and Teichner 1995; Krushkal and Quinn 2000].

If Σ is a union of discs or spheres, Theorem 1.2 follows from [Freedman and Quinn 1990, Theorem 10.5(1)]. The latter theorem contained an error found by Stong [1994] (see Theorem 5.7), but this is not relevant to Theorem 1.2 because of the way we defined the Kervaire–Milnor invariant. It is, however, very relevant to how to compute the Kervaire–Milnor invariant, and Stong’s correction will be one of the ingredients in our results (see Section 1.1).

For an arbitrary Σ , one could try to prove Theorem 1.2 by using general position to embed the 0– and 1–handles of Σ (relative to $\partial\Sigma$) and removing a small open neighbourhood thereof from M . This gives a new connected 4–manifold M_0 with the same fundamental group as M , and only the 2–handles $\{h_i : (D^2, S^1) \looparrowright (M_0, \partial M_0)\}$, one for each component Σ_i of Σ , remain to be embedded. One then hopes to apply [Freedman and Quinn 1990, Theorem 10.5(1)] (ie Theorem 1.2 for Σ a union of discs) to these maps of 2–handles to produce the desired embedded surface. The original algebraically dual spheres $\{g_i\}$ for the $\{f_i\}$ perform the same role for the $\{h_i\}$ in M_0 . The intersection and self-intersection numbers λ and μ remain unchanged; hence, they also vanish for $\{h_i\}$. However, the Kervaire–Milnor invariant may behave differently. That is, it may become nonzero for the embedding problem for the discs $\{h_i\}$, whereas it was trivial for the original F . We show that this phenomenon can occur in Example 9.3. This difference stems from the fact that, in applying [Freedman and Quinn 1990, Theorem 10.5(1)], we fix an embedding of the 1–skeleton of Σ and try to extend it across the 2–handles. As usual in obstruction theory, it might be advantageous to go back and change the solution of the problem on the 1–skeleton.

1.1 Computing the Kervaire–Milnor invariant

The strength of Theorem 1.2 versus the above strategy using [Freedman and Quinn 1990, Theorem 10.5(1)] lies in our computation of the Kervaire–Milnor invariant for arbitrary compact surfaces. In the case of discs and spheres, Stong showed that the Kervaire–Milnor invariant vanishes in more situations than claimed by Freedman and Quinn, due to the ambiguity arising from sheet choices when pairing up double points by Whitney discs, when the associated fundamental group element has order two. As we recall in Theorem 5.7, Stong [1994] introduced the notion of an r –characteristic surface, short for $\mathbb{R}P^2$ –characteristic surface (Definition 5.5), to give a criterion, in terms of copies of $\mathbb{R}P^2$ immersed in the ambient manifold M , to decide whether the sheet changing move is viable. Combined with the work of Freedman and Quinn, this enabled the computation of the Kervaire–Milnor invariant, and therefore answered the embedding problem for every finite union of discs or spheres with algebraically dual spheres, in an ambient 4–manifold with good fundamental group (see Remark 5.8 for more details).

In order to compute the Kervaire–Milnor invariant $\text{km}(F)$ for general surfaces, we extend the notion of r –characteristic surfaces to a notion of b –characteristic surfaces, short for *band characteristic* (Definition 5.17), defined using bands (annuli and Möbius bands) immersed in M . The next theorem is a generalisation of Stong’s computation of $\text{km}(F)$ to arbitrary compact surfaces.

Definition 1.5 Let $F = \{f_i\}_{i=1}^m : (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1 with $\mu(F) = 0$. Choose a convenient collection $\mathcal{W} = \{W_l\}$ of Whitney discs that pair all double points of F and define

$$t(F, \mathcal{W}) := \sum_{l,i} |\text{Int } W_l \pitchfork f_i| \pmod{2}.$$

In other words, $t(F, \mathcal{W})$ is the mod 2 count of transverse intersections between F and the interiors of the Whitney discs in \mathcal{W} .

We will often apply this definition to the restriction F° of $F = \{f_1, \dots, f_m\}$ to the subsurface $\Sigma^\circ \subseteq \Sigma$, which includes a component Σ_i of Σ precisely if its image does not admit a framed immersion $g_i: S^2 \looparrowright M$ with $\lambda(f_j, g_i) = \delta_{ij}$ for all $j = 1, \dots, m$ (Definition 5.1). The main result of the article is as follows.

Theorem 1.6 *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. Suppose that $\mu(F) = 0$ and that F has algebraically dual spheres. If F° is not b -characteristic then $\text{km}(F) = 0$. If F° is b -characteristic, then the secondary embedding obstruction satisfies*

$$\text{km}(F) = t(F^\circ, \mathcal{W}^\circ) \in \mathbb{Z}/2$$

for every convenient collection of Whitney discs \mathcal{W}° pairing all the double points of F° .

The main novelty in this theorem lies in finding the right condition on F that makes the combinatorial formula $t(F^\circ, \mathcal{W}^\circ)$ independent of the choice of Whitney discs, namely that F° is b -characteristic. Note that, if $\pi_1(M)$ is good, then, for $\text{km}(F) = 0$ in the previous theorem, Theorem 1.2 gives an embedding regularly homotopic to F . In practice, it can often be easy to determine if a given surface is b -characteristic, as demonstrated by the following corollaries, derived in Section 9 as consequences of the more general Proposition 9.1.

Corollary 1.7 *If M is a simply connected 4-manifold and Σ is a connected, oriented surface with positive genus, then any generic immersion $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ with vanishing self-intersection number is not b -characteristic. Thus, if F has an algebraically dual sphere, then $\text{km}(F) = 0$, and, since $\pi_1(M)$ is good, the map F is regularly homotopic, relative to $\partial\Sigma$, to an embedding.*

This corollary in particular implies that, for every simply connected 4-manifold M with boundary a disjoint union of homology spheres, every primitive class in $H_2(M; \mathbb{Z})$ can be represented by an embedded torus. This recovers [Lee and Wilczyński 1997, Theorem 1.1] in the case of divisibility $d = 1$. We also have the following extension to the case of arbitrary 4-manifolds.

Corollary 1.8 *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1, with $\mu(F) = 0$ and Σ connected. If F' is obtained from F by an ambient connected sum with an embedding $S^1 \times S^1 \hookrightarrow S^4$, then F' is not b -characteristic. Thus, if F has an algebraically dual sphere, then $\text{km}(F') = 0$, and if $\pi_1(M)$ is good, then F' is regularly homotopic, relative to $\partial\Sigma$, to an embedding.*

See Corollaries 1.13 and 1.14 for the nonorientable analogues of these two results. In particular, the latter concerns the case where we replace the embedded torus in Corollary 1.8 by an embedded $\mathbb{R}P^2$.

1.2 Band characteristic maps

We briefly explain how the notion of a map being b -characteristic arises in the context of embedding general surfaces. Given $F: \Sigma \looparrowright M$ as in Convention 1.1, assume that its double points are paired by a convenient collection of Whitney discs \mathcal{W} . Then the interiors of the discs in \mathcal{W} could be tubed into spheres




Figure 1: Two portions of the immersion F and part of a band B are shown on the left. A finger move produces F' with two new double points, paired by W_B .

in M , potentially changing the count t from Theorem 1.6. The condition that F is s -characteristic, short for *spherically characteristic* (Definition 5.2), precisely ensures that the count is preserved under this move.

Similarly, consider a band, ie an annulus or Möbius band, immersed in M with boundary lying on $F(\Sigma)$ minus the double points, as in Figure 1. Then, as shown in the figure, we may perform a finger move on F along a fibre of the band, creating F' with two new intersections, paired by a new Whitney disc W_B arising from the band B (see Figure 1). We call this move the *band fibre finger move* and give further details in Construction 7.2. Adding W_B to ${}^{\circ}W$ might in principle change the count t , but the requirement that F° is b -characteristic ensures it does not. In the case that Σ has only simply connected components, the boundary of the band is null-homotopic in Σ , and therefore the band can be closed off by discs to produce either a sphere (from an annulus) or an $\mathbb{R}P^2$ (from a Möbius band). Thus in this case it suffices to consider r -characteristic maps.

However, for general Σ there may exist a band in M with a boundary curve that is nontrivial in $\pi_1(\Sigma)$. This necessitates the new notion of b -characteristic maps, which by definition requires that a function $\Theta: \mathcal{B}(F) \rightarrow \mathbb{Z}/2$ vanishes (Definitions 5.12 and 5.14), where $\mathcal{B}(F)$ consists of the homology classes in $H_2(M, \Sigma; \mathbb{Z}/2)$ that can be represented by certain immersed bands in M with boundary on Σ (Definition 5.9). These additional conditions on the bands have to do with the first Stiefel–Whitney classes of M and Σ ; when both are oriented, $\mathcal{B}(F)$ consists precisely of the classes in $H_2(M, \Sigma; \mathbb{Z}/2)$ that are represented by maps of annuli and Möbius bands. Roughly speaking, the vanishing of Θ means that every band with boundary on Σ intersects Σ evenly many times in its interior. Intersections among the boundary components of the bands and a relative Euler number also play a role; see Sections 5 and 3.7 for details. If $\Theta \equiv 0$, then, for every band B , adding W_B does not change the t -count, and in fact t is well defined if and only if F is b -characteristic (Lemma 7.3). See Example 9.3 for a map which is r -characteristic but not b -characteristic.

The first step for deciding whether F is b -characteristic is to determine the subset $\mathcal{B}(F)$. In general this could be difficult, but in practice it is often soluble. If this can be done, then, as shown in Figure 2, by computing $\lambda_{\Sigma}|_{\partial\mathcal{B}(F)}$ and $\Theta: \mathcal{B}(F) \rightarrow \mathbb{Z}/2$, we can decide whether F is b -characteristic. Both of these are functions on a finite group, so in principle these computations are manageable.

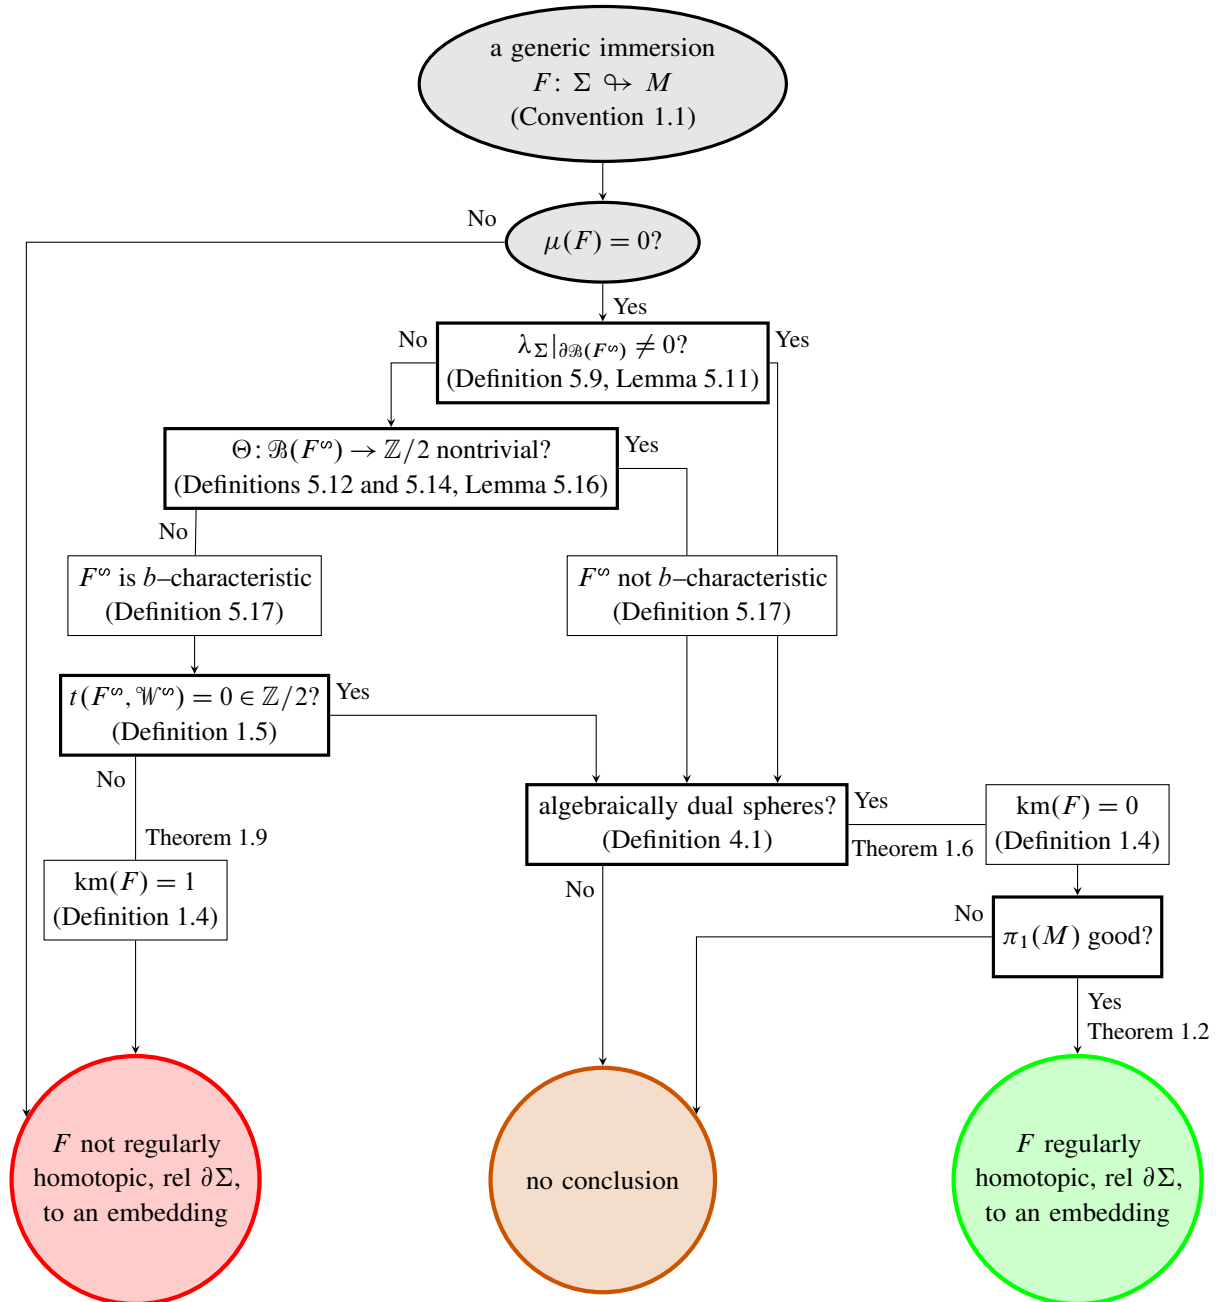


Figure 2: A flowchart deciding whether a generic immersion F is regularly homotopic, relative to the boundary, to an embedding.

1.3 An embedding obstruction without dual spheres

Irrespective of whether F has algebraically dual spheres, we obtain a secondary embedding obstruction in the b -characteristic case.

Theorem 1.9 *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1 with $\mu(F) = 0$. Let \mathcal{W} be a convenient collection of Whitney discs for the double points of F . Then F is b -characteristic if and only if, for every F' regularly homotopic to F and convenient collection \mathcal{W}' for the double points of F' , we have $t(F, \mathcal{W}) = t(F', \mathcal{W}')$.*

For b -characteristic F , we denote the resulting regular homotopy invariant by $t(F) \in \mathbb{Z}/2$. Then, if $\text{km}(F) = 0$ — for instance, if F is an embedding — then $t(F) = 0$.

Note that, in particular, if F is b -characteristic and a map H is regularly homotopic to F , then H is b -characteristic (Lemma 5.19). If F is not b -characteristic, it is still possible that some restriction F' of F to a union of connected components $\Sigma' \subseteq \Sigma$ is b -characteristic. Then we obtain an obstruction to embedding F' , and as a consequence to embedding F . A frequent example of this phenomenon is $F' = F^\infty$ from Theorem 1.6. Note that, by Lemma 5.3, if F is b -characteristic then $F = F^\infty$.

As part of our analysis of the obstruction t , in Section 9 we shall prove the following additivity properties.

Proposition 1.10 *Let M_1 and M_2 be oriented 4-manifolds. Let $F_1: (\Sigma_1, \partial\Sigma_1) \looparrowright (M_1, \partial M_1)$ and $F_2: (\Sigma_2, \partial\Sigma_2) \looparrowright (M_2, \partial M_2)$ be generic immersions of connected, compact, oriented surfaces, each with vanishing self-intersection number. If F_i is b -characteristic for each i , then both the disjoint union*

$$F_1 \sqcup F_2: \Sigma_1 \sqcup \Sigma_2 \looparrowright M_1 \# M_2$$

and any interior connected sum

$$F_1 \# F_2: \Sigma_1 \# \Sigma_2 \looparrowright M_1 \# M_2$$

are b -characteristic, and satisfy

$$t(F_1 \sqcup F_2) = t(F_1 \# F_2) = t(F_1) + t(F_2).$$

Theorem 1.9 and Proposition 1.10 imply the following corollary.

Corollary 1.11 *For any g , there exists a smooth, closed 4-manifold M_g , a closed, connected, oriented surface Σ_g of genus g , and a smooth, b -characteristic, generic immersion $F: \Sigma_g \looparrowright M_g$ with $t(F) \neq 0$, and therefore $\text{km}(F) \neq 0$.*

By contrast, we show in Example 9.5 that every map of a closed surface to $S^1 \times S^3$ is homotopic to an embedding. One could ask whether there exists a 4-manifold M and immersions $\Sigma_g \looparrowright M$ with nontrivial Kervaire–Milnor invariant, for every g . However, as a partial negative answer we will show in Proposition 9.7 that a b -characteristic generic immersion $F: \Sigma \looparrowright M$ from a closed surface Σ to a compact 4-manifold M with abelian fundamental group with n generators must satisfy $\chi(\Sigma) \geq -2n$.

1.4 Homotopy versus regular homotopy

Theorems 1.2, 1.9 and 1.6 together give a framework for deciding whether or not an immersed surface is regularly homotopic to an embedding, as illustrated by the flowchart in Figure 2. However, in the first

sentence of the article, we began by considering whether a given continuous map is homotopic to an embedding. We explain now how to extend the framework of Figure 2 to decide this, for maps of surfaces that admit algebraically dual spheres.

For a map f from a connected surface to a 4-manifold, we will show in Theorem 2.32 that there are either infinitely many or precisely two regular homotopy classes in the homotopy class of f , according to whether $f^*(w_1(M))$ is trivial or nontrivial, respectively. Our strategy is to make a judicious choice of regular homotopy class to which we apply our previous theory.

Begin with a continuous map $F: \Sigma \rightarrow M$ that restricts to an embedding on $\partial\Sigma$ and satisfies $F^{-1}(\partial M) = \partial\Sigma$, where Σ and M are as in Convention 1.1. Denote the components of F by $f_i: (\Sigma_i, \partial\Sigma_i) \rightarrow (M, \partial M)$, and suppose that F has algebraically dual spheres. Note that homotopies preserve the intersection numbers $\lambda(f_i, f_j)$, but might not preserve the self-intersection number $\mu(f_i)$, since adding a local cusp in f_i changes $\mu(f_i)_1$, the coefficient of $1 \in \pi_1(M)$, by ± 1 . Depending on the behaviour of the orientation characters of M and Σ , the coefficient $\mu(f_i)_1$ lies in either \mathbb{Z} or $\mathbb{Z}/2$, and is preserved under regular homotopy (see Lemma 2.24 and Proposition 2.25).

Now, in order to decide whether F is homotopic to an embedding, we will either find a generic immersion in the homotopy class of F which is regularly homotopic to an embedding, or show that this is impossible. First, by performing a homotopy we may assume without loss of generality that F is a generic immersion such that $\mu(f_i)_1 = 0$ for every component f_i of F . If $\mu(F) \neq 0$, then F is not homotopic to an embedding. On the other hand, if $\mu(F) = 0$, we have the two following cases. Below, $(f_i)_\bullet$ is the map induced on fundamental groups by f_i using some choice of path connecting $f_i(\Sigma_i)$ to the basepoint of M .

Case 1 $w_1(\Sigma)|_{\ker(f_i)_\bullet}$ is trivial for every $f_i \in F^\infty$.

By Theorem 2.32, the regular homotopy class of F^∞ is uniquely determined by the condition that $\mu(f_i)_1 = 0$ for each i with $f_i \in F^\infty$. Run the analysis in Figure 2 on F to determine whether it is regularly homotopic to an embedding. Note that the outcome of this analysis only depends on the regular homotopy class of F^∞ , rather than all of F . In particular, if $\pi_1(M)$ is good, then F is homotopic to an embedding if and only if F is regularly homotopic to an embedding.

Case 2 There exists $f_i \in F^\infty$ with $w_1(\Sigma)|_{\ker(f_i)_\bullet}$ nontrivial.

In this case we use the following theorem, which we prove in Section 6.

Theorem 1.12 *Let $F = \{f_i\}_{i=1}^m: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1 with $\mu(F) = 0$. Suppose that there is at least one $f_i \in F^\infty$ with $w_1(\Sigma)|_{\ker(f_i)_\bullet}$ nontrivial. Then there exists a generic immersion F' homotopic to F with $\mu(F') = 0$, and a convenient collection of Whitney discs ${}^{\circ}W'$ such that $t((F')^\infty, ({}^{\circ}W')^\infty) = 0$. Thus, if F' has algebraically dual spheres, then $\text{km}(F') = 0$, and if moreover $\pi_1(M)$ is good, then F' is regularly homotopic, relative to $\partial\Sigma$, to an embedding.*

Using this theorem, we can immediately conclude that our F as in Case 2 is homotopic to an embedding, as long as $\pi_1(M)$ is good. Notably, it is not relevant in this case whether F^∞ is b -characteristic. This completes the analysis of whether a given continuous map of a surface into a 4-manifold is homotopic to an embedding.

We now sketch the proof of Theorem 1.12. By the vanishing of $\mu(F)$, there is a convenient collection of Whitney discs ${}^{\circ}W$ for F and therefore for F^∞ . When $t(F^\infty, {}^{\circ}W^\infty) = 0$, the proof is completed by setting $F' = F$. When $t(F^\infty, {}^{\circ}W^\infty) = 1$, we use Construction 6.1 to find another generic immersion F' homotopic to F . Briefly, Construction 6.1 involves creating four new double points in the component f_i with nontrivial $w_1(\Sigma)|_{\ker(f_i)}$, using local cusps, and then cancelling them using a suitable choice of Whitney arcs and discs. For further details on the proof, see Section 6.

Theorem 1.12 also has the following immediate corollaries. These are the nonorientable analogues of Corollaries 1.7 and 1.8. They provide homotopies to embeddings rather than regular homotopies, and again it is not relevant whether F^∞ is b -characteristic.

Corollary 1.13 *If M is a simply connected 4-manifold and Σ is a connected, nonorientable surface, then a generic immersion $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ with vanishing self-intersection number and an algebraically dual sphere is homotopic, relative to $\partial\Sigma$, to an embedding.*

Corollary 1.14 *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1, with Σ connected and $\pi_1(M)$ good. Suppose that F has vanishing self-intersection number and an algebraically dual sphere. If F' is obtained from F by an ambient connected sum with any embedding $\mathbb{R}P^2 \hookrightarrow S^4$, then F' is homotopic, relative to $\partial\Sigma$, to an embedding.*

As in our analysis for the embedding problem up to regular homotopy, our techniques are primarily applicable in the presence of algebraically dual spheres and good fundamental group of the ambient space. It is however sometimes possible to conclude that a map without algebraically dual spheres is homotopic to an embedding. For example, we show in Example 9.5 that every map of a closed surface to $S^1 \times S^3$ is homotopic to an embedding.

1.5 Applications to knot theory

Theorem 1.2 can be applied to the problem of finding embedded surfaces in general 4-manifolds bounded by knots in their boundary. Given a closed 4-manifold M , let M° denote the punctured manifold $M \setminus \overset{\circ}{D}^4$. The M -genus of a knot $K \subseteq S^3 = \partial M^\circ$, denoted by $g_M(K)$, is the minimal genus of an embedded orientable surface bounding K in M° . If M is smooth, we also consider the *smooth M -genus*, denoted by $g_M^{\text{Diff}}(K)$, the minimal genus of a smoothly embedded orientable surface with boundary K . The quantities g_{S^4} and $g_{S^4}^{\text{Diff}}$ coincide with the topological and smooth slice genus of knots in \underline{D}^4 , respectively. Note that $g_{\overline{M}}(K) = g_M(\overline{K})$, so (2) and (3) below imply the corresponding results for $\mathbb{C}P^2$ and $*\mathbb{C}P^2$, respectively.

Corollary 1.15 For every knot $K \subseteq S^3$,

- (1) $g_M(K) = 0$ for every simply connected 4-manifold M not homeomorphic to one of S^4 , $\mathbb{C}\mathbb{P}^2$ or $*\mathbb{C}\mathbb{P}^2$;
- (2) $g_{\mathbb{C}\mathbb{P}^2}(K) \leq 1$ and $g_{\mathbb{C}\mathbb{P}^2}(\#^3 T(2, 3)) = 1$; and
- (3) $g_{*\mathbb{C}\mathbb{P}^2}(K) \leq 1$ and $g_{*\mathbb{C}\mathbb{P}^2}(\#^2 T(2, 3)) = 1$.

See Section 9 for the proof. The smooth $\mathbb{C}\mathbb{P}^2$ -genus has been studied by [Yasuhara 1991; 1992; Ait Nouh 2009; Pichelmeyer 2020; Marengon et al. 2024], and differs dramatically from the topological result in Corollary 1.15(2); in particular, it can be arbitrarily high [Marengon et al. 2024].

Corollary 1.15(1) is straightforward to prove if M topologically splits as a connected sum with $S^2 \times S^2$ or $S^2 \tilde{\times} S^2$, because $g_{S^2 \times S^2}(K) = g_{S^2 \tilde{\times} S^2}(K) = 0$ for all K by the Norman trick [1969, Corollary 3, Remark]. For the $K3$ surface, this implies that $g_{K3}(K) = 0$ for all knots K . On the other hand, it is an open question whether there exists a K with $g_{K3}^{\text{Diff}}(K) \neq 0$ [Manolescu et al. 2024, Question 6.1].

Given a knot $K \subseteq S^3$ and an integer $n \in \mathbb{Z}$, we build the n -trace $X_n(K)$ by attaching a 2-handle D^4 along K with framing n . The minimal genus of an embedded surface representing a generator of $H_2(X_n(K); \mathbb{Z})$ is called the n -shake genus of K , denoted by $g_n^{\text{sh}}(K)$. Similarly, the smooth n -shake genus of K is denoted by $g_n^{\text{sh,Diff}}(K)$. We recover the following result of [Feller et al. 2021].

Corollary 1.16 [Feller et al. 2021, Proposition 8.7] For any knot $K \subseteq S^3$, $g_{\pm 1}^{\text{sh}}(K) = \text{Arf}(K) \in \{0, 1\}$.

By contrast, the smooth ± 1 -shake genus of a knot can be arbitrarily high. For example, for $q \geq 5$ we have that $g_{\pm 1}^{\text{sh,Diff}}(T(2, q)) \geq \frac{1}{2}(q + 1)$, by the slice-Bennequin inequality [Lisca and Matić 1998; Cochran and Ray 2016, Corollary 5.2].

Outline of the paper

In Section 2, we describe the primary embedding obstructions arising from the theory of equivariant intersection numbers for surfaces in 4-manifolds. In Section 3, we define the Kervaire–Milnor invariant carefully. We prove Theorem 1.2 in Section 4. In Section 5, we explain our combinatorial method for computing the Kervaire–Milnor invariant, and define b -characteristic surfaces, postponing almost all proofs to Section 7. In Section 6, we give the proof of Theorem 1.12. In Section 8, we prove Theorems 1.6 and 1.9. Finally, in Section 9, we prove Corollaries 1.7, 1.8, 1.11, 1.15 and 1.16 and Proposition 1.10, and we give further applications and examples.

Acknowledgements

We are grateful to Rob Schneiderman and the referee for several insightful comments on a previous version, and to Allison N Miller and Andrew Nicas for helpful conversations. Much of this research was conducted at the Max Planck Institute for Mathematics. Kasprowski was supported by the Deutsche

Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy — GZ 2047/1, Projekt-ID 390685813. Powell was partially supported by EPSRC New Investigator grant EP/T028335/2 and EPSRC New Horizons grant EP/V04821X/2.

2 Generic immersions and intersection numbers

In Section 2.1, we carefully define and study generic immersions of surfaces in 4–manifolds in the topological category. We show they admit well-behaved normal bundles, and introduce generic homotopies and ambient isotopies between them.

In Sections 2.2 and 2.3, we study equivariant intersection and self-intersection numbers of generically immersed surfaces. In the case of immersions of spheres and discs, these have a long history, in particular in surgery theory (see eg [Wall 1970]). For the first time in the literature, as far as we are aware, we give a careful account of intersection and self-intersection numbers in full generality for compact surfaces and for any possible combination of orientation characters. The specific groups in which these numbers live depend on the input surfaces, and it is somewhat subtle to describe them. A preliminary version for orientable surfaces was considered in eg [Cochran et al. 2003, Section 7], and the self-intersection number for annuli was considered in [Schneiderman 2003].

In Section 2.4, we discuss Whitney discs, which arise if the intersection and self-intersection numbers vanish. We define the important notion of a convenient collection of Whitney discs. In Section 2.5, Theorem 2.32 explains the difference between homotopy and regular homotopy of generic immersions of surfaces, in terms of the Euler number of the normal bundle or the self-intersection number.

2.1 Topological generic immersions

We start with the definition of an immersion of manifolds in the topological setting. For $m \geq 0$, let $\mathbb{R}_+^m := \{(x_1, \dots, x_m) \in \mathbb{R}^m \mid x_1 \geq 0\}$. For $k \leq n$, we consider the standard inclusions

$$\begin{aligned} \iota: \mathbb{R}^k &= \mathbb{R}^k \times \{0\} \hookrightarrow \mathbb{R}^k \times \mathbb{R}^{n-k} = \mathbb{R}^n, \\ \iota_+: \mathbb{R}_+^k &= \mathbb{R}_+^k \times \{0\} \hookrightarrow \mathbb{R}^k \times \mathbb{R}^{n-k} = \mathbb{R}^n, \\ \iota_{++}: \mathbb{R}_+^k &= \mathbb{R}_+^k \times \{0\} \hookrightarrow \mathbb{R}_+^k \times \mathbb{R}^{n-k} = \mathbb{R}_+^n. \end{aligned}$$

Definition 2.1 A continuous map $F: \Sigma^k \rightarrow M^n$ between topological manifolds of dimensions $k \leq n$ is an *immersion* if locally it is a flat embedding, that is, if, for each point $p \in \Sigma$, there is a chart φ around p and a chart Ψ around $F(p)$ fitting into one of the commutative diagrams

$$(2-1) \quad \begin{array}{ccc} \mathbb{R}^k & \xrightarrow{\iota} & \mathbb{R}^n \\ \downarrow \varphi & & \downarrow \Psi \\ \Sigma & \xrightarrow{F} & M \end{array} \quad \begin{array}{ccc} \mathbb{R}_+^k & \xrightarrow{\iota_+} & \mathbb{R}^n \\ \downarrow \varphi & & \downarrow \Psi \\ \Sigma & \xrightarrow{F} & M \end{array} \quad \begin{array}{ccc} \mathbb{R}_+^k & \xrightarrow{\iota_{++}} & \mathbb{R}_+^n \\ \downarrow \varphi & & \downarrow \Psi \\ \Sigma & \xrightarrow{F} & M \end{array}$$

The first diagram is for $p \in \text{Int } \Sigma$ and $F(p) \in \text{Int } M$, the second diagram is for $p \in \partial \Sigma$ and $F(p) \in \text{Int } M$, and the third is for $p \in \partial \Sigma$ and $F(p) \in \partial M$. In particular, F is required to map interior points of Σ to interior points of M .

Some authors prefer to call this notion a *locally flat immersion*.

Definition 2.2 A (linear) *normal bundle* for an immersion $F: \Sigma^k \rightarrow M^n$ is an $(n-k)$ -dimensional real vector bundle $\pi: \nu_F \rightarrow \Sigma$, together with an immersion $\tilde{F}: \nu_F \rightarrow M$ that restricts to F on the zero section s_0 , ie $\tilde{F} \circ s_0 = F$, and such that each point $p \in \Sigma$ has a neighbourhood U such that $\tilde{F}|_{\pi^{-1}(U)}$ is an embedding.

We now restrict to the relevant dimensions for this paper, $k = 2$ and $n = 4$, and take M to be a connected topological 4-manifold as in Convention 1.1. The *singular set* of an immersion $F: \Sigma \rightarrow M$ is the set

$$\mathcal{S}(F) := \{m \in M : |F^{-1}(m)| \geq 2\}.$$

Recall that a continuous map is said to be *proper* if the inverse image of every compact set in the codomain is compact.

Definition 2.3 Let Σ be a surface, possibly noncompact. A continuous, proper map $F: \Sigma \rightarrow M$ is said to be a (topological) *generic immersion*, denoted by $F: \Sigma \looparrowright M$, if it is an immersion and the singular set is a closed, discrete subset of M consisting only of transverse double points, each of whose preimages lies in the interior of Σ . In particular, whenever $m \in \mathcal{S}(F)$, there are exactly two points $p_1, p_2 \in \Sigma$ with $F(p_i) = m$, and there are disjoint charts φ_i around p_i for $i = 1, 2$, where φ_1 is as in the leftmost diagram of (2-1) and φ_2 is the same, with respect to the same chart Ψ around m , but with ι replaced by

$$\iota': \mathbb{R}^2 = \{0\} \times \mathbb{R}^2 \hookrightarrow \mathbb{R}^2 \times \mathbb{R}^2 = \mathbb{R}^4.$$

Theorem 2.4 A generic immersion $F: \Sigma \looparrowright M$, for possibly noncompact Σ , has a normal bundle as in Definition 2.2 with the additional property that \tilde{F} is an embedding outside a neighbourhood of $F^{-1}(\mathcal{S}(F))$, and near the double points \tilde{F} plumbs two coordinate regions $\pi^{-1}(\varphi_i(\mathbb{R}^2)) \cong \varphi_i(\mathbb{R}^2) \times \mathbb{R}^2$ for $i = 1, 2$ together, ie $\tilde{F} \circ (\varphi_1(x), y) = \tilde{F} \circ (\varphi_2(y), x)$.

Proof Let $\partial_1 \Sigma \subseteq \partial \Sigma$ denote the union of the components of $\partial \Sigma$ mapped to ∂M . Then, since $F|_{\partial_1 \Sigma}$ is an embedding of a 1-manifold in a 3-manifold, it has a normal bundle. We extend this to a collar neighbourhood of $F(\partial_1 \Sigma)$ contained in a collar neighbourhood of ∂M . To do this, first note that, since ∂M is closed in M , $\mathcal{S}(F)$ is closed and contained in $\text{Int}(M)$, and manifolds are normal, it follows that there is an open neighbourhood of ∂M disjoint from $\mathcal{S}(F)$. Then argue as in Connelly's proof [1971] that boundaries of manifolds have collars, to obtain a homeomorphism of pairs

$$G: (M, F(\Sigma)) \xrightarrow{\cong} (M \cup (\partial M \times [0, 1]), F(\Sigma) \cup (F(\partial_1 \Sigma) \times [0, 1])).$$

The normal bundle over the boundary extends to a collar in the codomain; hence, its pullback extends to a collar in the domain.

Next, let $\partial_2 \Sigma \subseteq \partial \Sigma$ denote the union of the components of $\partial \Sigma$ mapped to $\text{Int } M$. We see that $F(\partial_2 \Sigma)$ has a normal bundle by [Freedman and Quinn 1990, Theorem 9.3], as a submanifold of M . Let \tilde{F} be the embedding of the total space, as in Definition 2.2. By using the inward-pointing normal for $\partial_2 \Sigma$ in Σ , we obtain an orthogonal decomposition of each fibre as $\nu_{\partial_2 \Sigma \hookrightarrow \Sigma} \oplus V$, where V is a 2-dimensional subspace. Then translates of V in the direction of the inward-pointing normal give rise to a normal bundle on the intersection of a collar of $\partial_2 \Sigma$ with the image of the normal bundle of $\partial_2 \Sigma$ under \tilde{F} .

Now we want to extend the normal bundle that we have just constructed on a neighbourhood of $\partial \Sigma$ to the rest of Σ . First we will produce a normal bundle in a neighbourhood of both preimages of each double point, and then finally we will extend the normal bundle to the rest of the interior of Σ .

Let $m \in \mathcal{S}(F)$ be a double point of F , so that there exist $p_1, p_2 \in \Sigma$ with $F(p_i) = m$ for $i = 1, 2$. By the definition of a generic immersion, there is a chart Ψ for M at m , and charts φ_i around p_i , such that $F \circ \varphi_1(x) = \Psi(x, 0)$ and $F \circ \varphi_2(y) = \Psi(0, y)$. We assume that $F(\Sigma) \cap \Psi(\mathbb{R}^4) = F(\varphi_1(\mathbb{R}^2)) \cup F(\varphi_2(\mathbb{R}^2))$, and moreover that the images of the charts for different elements of $\mathcal{S}(F)$ do not overlap one another, and also are disjoint from the images of the normal bundles already constructed close to $\partial \Sigma$. Then we take a trivial \mathbb{R}^2 -bundle over each $\varphi_i(\mathbb{R}^2)$, and we define the map \tilde{F} on $\varphi_1(\mathbb{R}^2) \times \mathbb{R}^2$ and $\varphi_2(\mathbb{R}^2) \times \mathbb{R}^2$ by setting $\tilde{F}(\varphi_1(x), y) = \Psi(x, y)$ and $\tilde{F}(\varphi_2(x), y) = \Psi(y, x)$. Then $\tilde{F} \circ (\varphi_2(y), x) = \Psi(x, y) = \tilde{F} \circ (\varphi_1(x), y)$, as needed.

Let U_1^m and U_2^m be open neighbourhoods in Σ of p_1 and p_2 , contained within the images of φ_1 and φ_2 above, respectively. Define $\Sigma' := \Sigma \setminus \bigcup_{m \in \mathcal{S}(F)} (U_1^m \cup U_2^m)$. Then the restriction of F gives an embedding of Σ' in M . We already have a normal bundle defined on a neighbourhood of $\partial \Sigma'$. Apply [Freedman and Quinn 1990, Theorem 9.3A] to extend the given normal bundle on $\overline{U_1^m \cup U_2^m}$ and $\partial \Sigma$ to all of Σ' and therefore we have a normal bundle on all of Σ . \square

Remark 2.5 Freedman and Quinn [1990, Theorem 9.3] produce an *extendable* normal bundle for every submanifold of a 4-manifold. The extendibility condition is technical with an important consequence: extendable normal bundles are unique up to isotopy. One can always find an extendable normal bundle embedded in the total space of any given normal bundle.

The proof of Theorem 2.4 also applies in the more general setting where $\partial_2 \Sigma$ is not embedded in M , but $F|_{\partial_2 \Sigma}$ factors as a composition of generic immersions $\partial_2 \Sigma \looparrowright S \looparrowright M$ for some surface S . We will use this case in the definition of b -characteristic maps in Section 5, so we introduce nomenclature.

Definition 2.6 Let $g: S \looparrowright M$ be a generic immersion of a surface in a 4-manifold M . Let (B, Z) be a pair consisting of a surface B and a collection $Z \subseteq \partial B$ of connected components of its boundary. A map $H: B \rightarrow M$ is called a *generic immersion of pairs* if $H(Z) \subseteq g(S)$ and

- (i) $H|_{B \setminus Z}$ is a generic immersion that is transverse to g and has image disjoint from $H(Z)$;
- (ii) $H(B)$ is disjoint from the double points of g , which implies there is a unique map $h: Z \rightarrow S$ with $g \circ h = H|_Z$;

- (iii) the map h is a generic immersion; and
- (iv) there is a collar N of Z in B with $H(N \setminus Z) \subseteq M \setminus g(S)$.

We denote such maps by $H: (B, Z) \looparrowright (M, S)$, and sometimes identify h with $H|_Z$.

Corollary 2.7 *Let $g: S \looparrowright M$ be a generic immersion of a surface in a 4-manifold M and let $H: (B, Z) \looparrowright (M, S)$ be a generic immersion of pairs. Then H admits a normal bundle, ie a normal bundle for B in M such that the restriction to Z contains a normal bundle for Z in S .*

Proof Note that Z has a normal bundle in S , and then the sum of this with the normal bundle of S in M guaranteed by Theorem 2.4 gives rise to a normal bundle for Z in M . The rest of the proof proceeds as before. □

Observe that smooth generic immersions are topological immersions. Next we show that when both notions make sense they coincide, which justifies the terminology.

Theorem 2.8 *Consider a smooth compact surface Σ and a (topological) generic immersion $F: \Sigma \looparrowright M$. If M is noncompact then let $M' := M$, and if M is compact then choose $p \in M \setminus F(\Sigma)$ and set $M' := M \setminus \{p\}$. Then F is a smooth generic immersion in some smooth structure on M' .*

We know that M' has a smooth structure by [Freedman and Quinn 1990, Theorem 8.2; Quinn 1982, Corollary 2.2.3].

Proof Fix a smooth structure on ∂M such that the generic immersion F restricted to those connected components of $\partial \Sigma$ that map to ∂M is a smooth embedding. To find such a smooth structure, first use the standard smooth structure on the normal bundle of $F|_{\partial \Sigma}$, and then extend this to a smooth structure on all of ∂M . Since any two smooth structures on a 3-manifold are isotopic, this could also be arranged by an isotopy of $\partial \Sigma$, but our aim is to use the given map without isotoping it.

By Theorem 2.4, there is a normal bundle (ν_F, \tilde{F}) for F . Let $D(\nu_F) \rightarrow \Sigma$ be the (closed) disc bundle. This yields a regular neighbourhood $N(F) := \tilde{F}(D(\nu_F))$ of $F(\Sigma)$, a codimension zero submanifold of M' . The regular neighbourhood $N(F)$ can be identified with a smooth manifold obtained from $D(\nu_F)$ after the requisite plumbing operations and smoothing corners. Use such an identification to fix a smooth structure on $N(F) \subseteq M$. With respect to this smooth structure, the map $\Sigma \rightarrow N(F)$ is a smooth generic immersion.

Now, the boundary of $N(F)$ inherits a smooth structure. The complement of $\text{Int } N(F) \cup (N(F) \cap \partial M')$ in M' is a connected, noncompact 4-manifold with a prescribed smooth structure on its boundary. Then the interior has a compatible smooth structure by [Freedman and Quinn 1990, Theorem 8.2; Quinn 1982, Corollary 2.2.3], giving rise to a smooth structure on all of M' . Since the smooth structure on $N(F)$ is unaltered, F become a smooth generic immersion, as desired. □

Recall that an *isotopy* of homeomorphisms of a manifold M is a map $H: M \times [0, 1] \rightarrow M$ such that the track $M \times [0, 1] \rightarrow M \times [0, 1]$ given by $(m, t) \mapsto (H(m, t), t)$ is a homeomorphism.

Definition 2.9 An *ambient isotopy* between generic immersions $F, G: \Sigma \looparrowright M$ consists of two isotopies $H_\Sigma: \Sigma \times [0, 1] \rightarrow \Sigma$ and $H_M: M \times [0, 1] \rightarrow M$ such that

- (1) $H_\Sigma(-, 0)$ and $H_M(-, 0)$ are both the identity; and
- (2) $G(x) = H_M(F(H_\Sigma(x, 1)), 1)$ for all $x \in \Sigma$.

This is motivated by the smooth result which states that two generic immersions are ambiently isotopic (in the sense of Definition 2.9 but with homeomorphism replaced by diffeomorphism in the definition of an isotopy) if and only if they are connected by a path in the space of generic immersions [Golubitsky and Guillemin 1973, Chapter III, Theorem 3.11]. Note that for embeddings one does not need the isotopy H_Σ .

Mirroring the smooth notion, a *generic homotopy* between generically immersed surfaces in a 4–manifold is by definition a sequence of ambient isotopies, finger moves, Whitney moves and cusp homotopies. The moves in question are defined in local coordinates exactly as in the smooth setting. A *regular homotopy* between generically immersed surfaces in a 4–manifold is by definition a sequence of ambient isotopies, finger moves and Whitney moves. The following proposition explains that maps of surfaces in a 4–manifold can be assumed to be generic immersions, and homotopies between generic immersions may be assumed to be generic as well.

Proposition 2.10 [Powell et al. 2020, Proposition 3.1] *Let Σ be a compact surface and let M be a topological 4–manifold.*

- (1) *Every map $(\Sigma, \partial\Sigma) \rightarrow (M, \partial M)$ is homotopic (relative to the embedded boundary) to a generic immersion.*
- (2) *Every homotopy $(\Sigma, \partial\Sigma) \times [0, 1] \rightarrow (M, \partial M)$ that restricts to a generic immersion on $\Sigma \times \{0, 1\}$ is homotopic (relative to the boundary) to a generic homotopy.*

Briefly, the proposition is proven as follows. Homotope the maps away from a point of M using cellular approximation, remove that point, choose a smooth structure on the complement of the point, and then apply the smooth theory of generic immersions, combining [Hirsch 1976, Theorems 2.2.6 and 2.2.12] with [Golubitsky and Guillemin 1973, Chapter III, Corollary 3.3].

2.2 Intersection numbers

We define intersection numbers between compact, connected surfaces in 4–manifolds. In order to accommodate the fundamental group in equivariant intersection numbers, we need to use basings.

Definition 2.11 We call a manifold X *based* if it is equipped with basepoints $p_i \in X_i$ for each connected component $X_i \subseteq X$, together with a local orientation at each p_i . A generic immersion $F: X \rightarrow Y$ between based manifolds with Y connected is said to be *based* if it is equipped with *whiskers*, ie paths in Y from the basepoint of Y to $F(p_i)$ for each basepoint p_i of X .

For the remainder of this section, let M be a connected, based 4-manifold and let Σ and Σ' be based, compact, connected surfaces, unless specified otherwise.

Let $f: \Sigma \rightarrow M$ and $g: \Sigma' \rightarrow M$ be based maps that are *transverse*, ie around each intersection point $f(s) = g(s')$ with $s \in \Sigma$ and $s' \in \Sigma'$, there are coordinates that make f and g (in a neighbourhood of s in Σ and a neighbourhood of s' in Σ') resemble the standard inclusions $\mathbb{R}^2 \times \{0\}$ and $\{0\} \times \mathbb{R}^2$ into \mathbb{R}^4 , respectively, as in (2-1). We assume that these intersections are the only singularities between f and g and that $f(\partial\Sigma)$ and $g(\partial\Sigma')$ are disjoint.

Let v_f and v_g be whiskers for f and g . The intersection number $\lambda(f, g)$ is the sum of signed fundamental group elements

$$\lambda(f, g) := \sum_{p \in f \pitchfork g} \varepsilon(p) \cdot \eta(p)$$

as follows. A priori this is the formal sum of a list of elements of the set $\{\pm 1\} \times \pi_1(M)$. It will ultimately give rise to an element of a quotient of $\mathbb{Z}[\pi_1(M)]$, given in Definition 2.12, after we factor out the effect of finger and Whitney moves and the effect of the choice of the paths γ_f^p and γ_g^p in the first bullet point below.

Fix $p \in f \pitchfork g$. Next we define $\varepsilon(p) \in \{\pm 1\}$ and $\eta(p) \in \pi_1(M)$. We use $*$ to denote concatenation of paths.

- Let γ_f^p be a path in Σ from the basepoint to $f^{-1}(p)$ and let γ_g^p be a path in Σ' from the basepoint to $g^{-1}(p)$.
- The sign $\varepsilon(p) \in \{\pm 1\}$ is determined as follows. Transport the local orientation of Σ at the basepoint to $f^{-1}(p)$ along γ_f^p , and the local orientation of Σ' at the basepoint to $g^{-1}(p)$ along γ_g^p . This induces a local orientation at p , by ordering f before g . Another local orientation is obtained by transporting the local orientation at the basepoint of M to p along the concatenated path $v_g * (g \circ \gamma_g^p)$. We define $\varepsilon(p) = +1$ when the two local orientations match at p , and -1 otherwise.
- The element $\eta(p) \in \pi_1(M)$ is by definition the concatenation $v_f * (f \circ \gamma_f^p) * (g \circ \gamma_g^p)^{-1} * v_g^{-1}$.

For a generic immersion $f: \Sigma \looparrowright M$, we define $\lambda(f, f) := \lambda(f, f^+)$, where f^+ is a push-off of f along a section of its normal bundle transverse to the zero section. If the embedding $f|_{\partial\Sigma}$ is equipped with a specified framing for its normal bundle, then f^+ is defined to be a push-off of f along a section restricting to the first vector of that framing on $\partial\Sigma$.

If f_t is a homotopy of f that is transverse to g for all t then $\lambda(f_t, g)$ is independent of t as a set of signed fundamental group elements, assuming the above choices of $\gamma_{f_t}^p$ are made carefully. However, if f_t describes a finger move of f into g , there is a single time t_0 at which f_{t_0} and g are not transverse, because there is a tangency. After the tangency, two new intersection points p and q arise. These have the same group element $\eta(p) = \eta(q)$ and opposite signs $\varepsilon(p) = -\varepsilon(q)$, with appropriate choices of $\gamma_{f_t}^p$ and $\gamma_{f_t}^q$. Similarly, a Whitney move reduces the intersections between f and g by such a pair. To get a regular homotopy invariant notion, it is thus important to specify the home of $\lambda(f, g)$ carefully.

For Σ and Σ' simply connected, the sum $\lambda(f, g)$ is usually considered as an element of $\mathbb{Z}[\pi_1(M)]$ and is independent of the choice of $\{\gamma_f^p\}_p$ and $\{\gamma_g^p\}_p$. In the abelian group $\mathbb{Z}[\pi_1(M)]$ the relations $-a + a = 0$ for each $a \in \pi_1(M)$ are built in, and if one identifies the sign $\varepsilon(p)$ with the inverse in this abelian group then finger moves and Whitney moves do not change $\lambda(f, g)$ as an element in the group ring.

For nonsimply connected Σ and Σ' , the homotopy class of γ_f^p and γ_g^p may be changed by wrapping around nontrivial elements in $\pi_1(\Sigma)$ or $\pi_1(\Sigma')$. This wrapping may also change the induced local orientations at the intersection points of f and g . We describe this in more detail next. Let $w^M : \pi_1(M) \rightarrow \{\pm 1\}$, $w^\Sigma : \pi_1(\Sigma) \rightarrow \{\pm 1\}$ and $w^{\Sigma'} : \pi_1(\Sigma') \rightarrow \{\pm 1\}$ denote the orientation characters.

Definition 2.12 Let $\Gamma_{f,g}$ be the abelian group generated by the elements of $\pi_1(M)$ and with relators

$$(2-2) \quad \gamma - w^\Sigma(\alpha)w^{\Sigma'}(\beta)w^M(g_\bullet(\beta)) \cdot f_\bullet(\alpha) * \gamma * g_\bullet(\beta),$$

for all $\alpha \in \pi_1(\Sigma)$, $\beta \in \pi_1(\Sigma')$, $\gamma \in \pi_1(M)$. Here $f_\bullet(\alpha) := v_f * (f \circ \alpha) * v_f^{-1}$ and $g_\bullet(\beta) := v_g * (g \circ \beta) * v_g^{-1}$ are elements of $\pi_1(M)$.

For transverse $f : \Sigma \rightarrow M$ and $g : \Sigma' \rightarrow M$, the intersection number $\lambda(f, g) \in \Gamma_{f,g}$ is well defined. The relations precisely account for wrapping around elements of $\pi_1(\Sigma)$ or $\pi_1(\Sigma')$ as described above. We will show in Proposition 2.18 that this target also makes $\lambda(f, g)$ a homotopy invariant.

Remark 2.13 In the case that M , Σ and Σ' are all oriented,

$$\Gamma_{f,g} \cong \mathbb{Z}[f_\bullet(\pi_1(\Sigma)) \backslash \pi_1(M) / g_\bullet(\pi_1(\Sigma'))],$$

the free abelian group generated by the double coset quotient of $\pi_1(M)$ by left and right multiplication by the images of loops in Σ and Σ' , respectively.

In general, due to the signs introduced by the orientation characters, there may be torsion in $\Gamma_{f,g}$. For example, consider $f : \mathbb{RP}^2 \looparrowright M$ with $f_\bullet(\mathbb{RP}^1) = 1$, without any assumption on M . Then, for every $g : \Sigma' \rightarrow M$, the group

$$\Gamma_{f,g} \cong (\mathbb{Z}/2)[\pi_1(M) / g_\bullet(\pi_1(\Sigma'))]$$

is 2-torsion, due to the relations $\gamma = w^{\mathbb{RP}^2}(\mathbb{RP}^1) \cdot f_\bullet(\mathbb{RP}^1) * \gamma = -\gamma$ for every $\gamma \in \pi_1(M)$, arising from setting $\alpha := \mathbb{RP}^1$ and $\beta := 1$ in (2-2).

To understand $\Gamma_{f,g}$ better, we introduce some notation. Write $\pm\pi_1(M) := \{\pm 1\} \times \pi_1(M)$. There is a natural inclusion $\pm\pi_1(M) \rightarrow \mathbb{Z}[\pi_1(M)]$ given by $(\pm 1, \gamma) \mapsto \pm\gamma$. Write $[a] \in \Gamma_{f,g}$ for the equivalence class of $a \in \mathbb{Z}[\pi_1(M)]$, and let \sim denote the equivalence relation on $\pm\pi_1(M)$ induced by the composition $\pm\pi_1(M) \hookrightarrow \mathbb{Z}[\pi_1(M)] \twoheadrightarrow \Gamma_{f,g}$, ie for $a, b \in \pm\pi_1(M)$, $a \sim b$ if and only if the images of a and b in $\Gamma_{f,g}$ coincide. The following lemma is immediate from the definitions.

Lemma 2.14 *Let $\gamma_1, \gamma_2 \in \pi_1(M)$. One of the relations $[\gamma_1] = \pm[\gamma_2] \in \Gamma_{f,g}$ holds if and only if γ_1 and γ_2 represent the same element in the double coset $f_\bullet(\pi_1(\Sigma)) \backslash \pi_1(M) / g_\bullet(\pi_1(\Sigma'))$.*

Let $\text{pm}: \pm\pi_1(M) / \sim \twoheadrightarrow f_\bullet\pi_1(\Sigma) \backslash \pi_1(M) / g_\bullet\pi_1(\Sigma')$ be the map sending $\pm\gamma$ to the class of γ . We write $|\gamma| := \text{pm}(\gamma)$. Here one should think that pm stands for dividing out “plus–minus”. Note that pm has fibres of order 1 or 2 and we can decompose the double coset as a disjoint union $B_1 \sqcup B_2$ according to this distinction, where pm gives a bijection $\text{pm}^{-1}(B_1) \leftrightarrow B_1$ while $\text{pm}^{-1}(B_2) \rightarrow B_2$ is two-to-one.

Remark 2.15 We give examples in the cases from Remark 2.13. In the case that M, Σ and Σ' are all oriented, $B_1 = \emptyset$ and $B_2 = f_\bullet(\pi_1(\Sigma)) \backslash \pi_1(M) / g_\bullet(\pi_1(\Sigma'))$. If we have $f: \mathbb{RP}^2 \looparrowright M$ with $f_\bullet(\mathbb{RP}^1) = 1$, and $g: \Sigma' \rightarrow M$ is arbitrary, then $B_1 = \pi_1(M) / g_\bullet(\pi_1(\Sigma'))$ and $B_2 = \emptyset$.

Choose a section s of pm . For each $s(b) \in \text{pm}^{-1}(B_2)$, we denote the other element of $\text{pm}^{-1}(b)$ by $-s(b)$. Their images in $\Gamma_{f,g}$ are indeed inverse to one another, which motivates the notation.

Lemma 2.16 *Fix a section s for pm as above. The abelian group $\Gamma_{f,g}$ is a direct sum $\Gamma_{f,g} = FA \oplus V$ of a free abelian group FA on the set $s(B_2) \subseteq \pm\pi_1(M) / \sim \subseteq \Gamma_{f,g}$ and a $\mathbb{Z}/2$ -vector space V with basis $s(B_1) = \text{pm}^{-1}(B_1) \subseteq \pm\pi_1(M) / \sim$.*

Reading off the coefficients in this decomposition gives homomorphisms $c_{s(b)}: \Gamma_{f,g} \rightarrow \mathbb{Z}$ for each $b \in B_2$, and $c_b: \Gamma_{f,g} \rightarrow \mathbb{Z}/2$ for each $b \in B_1$, yielding a decomposition of $\Gamma_{f,g}$ as a direct sum of copies of \mathbb{Z} and $\mathbb{Z}/2$.

In particular, the homomorphisms $c_{s(b)}$ and c_b determine the isomorphisms displayed in Remark 2.13.

Proof Starting with the free abelian group with basis $\pi_1(M)$, a relator in (2-2) does one of the following three things:

- (1) It identifies two distinct basis elements γ_1 and γ_2 if and only if $\gamma_2 = f_\bullet(\alpha) * \gamma_1 * g_\bullet(\beta) \in \pi_1(M)$ and $w^\Sigma(\alpha)w^{\Sigma'}(\beta)w^M(g_\bullet(\beta)) = 1$ for some $\alpha \in \pi_1(\Sigma)$ and $\beta \in \pi_1(\Sigma')$.
- (2) It identifies a basis element γ_1 with the inverse $-\gamma_2$ of another basis element $\gamma_2 \neq \gamma_1$ if and only if $\gamma_2 = f_\bullet(\alpha) * \gamma_1 * g_\bullet(\beta)$ and $w^\Sigma(\alpha)w^{\Sigma'}(\beta)w^M(g_\bullet(\beta)) = -1$ for some $\alpha \in \pi_1(\Sigma)$ and $\beta \in \pi_1(\Sigma')$.
- (3) It identifies a basis element γ with its inverse $-\gamma$ if and only if

$$\gamma = f_\bullet(\alpha) * \gamma * g_\bullet(\beta) \quad \text{and} \quad w^\Sigma(\alpha)w^{\Sigma'}(\beta)w^M(g_\bullet(\beta)) = -1$$

for some $\alpha \in \pi_1(\Sigma)$ and $\beta \in \pi_1(\Sigma')$.

The first two types of relators reduce the basis to the double coset

$$f \bullet \pi_1(\Sigma) \backslash \pi_1(M) / g \bullet \pi_1(\Sigma').$$

The third type adds the relations $2[\gamma] = 0$ to the fundamental group elements γ in question, which then generate V , because these are exactly the γ such that $-[\gamma] = [\gamma]$, ie where $\# \text{pm}^{-1}(|\gamma|) = 1$. Those γ where $\# \text{pm}^{-1}(|\gamma|) = 2$ remain of infinite order and generate FA . Note that the second type of relator forces us to choose the section s in order to write down a consistent basis for FA . \square

The subgroup V and its basis clearly do not depend on our choice of section, but the basis of FA depends on this choice. If we change the section s at a point $b \in B_2$ to s' so that $s'(b) = -s(b)$, the associated basis element changes to its inverse. It follows that the subgroup FA of $\Gamma_{f,g}$ does not depend on the choice of s . It also follows that the coefficient maps $c_s(b)$ only depend on s up to sign and satisfy $c_{-s(b)} = -c_s(b)$.

For a given $a \in \pm \pi_1(M) / \sim$, we can choose a section s as above with $s(\text{pm}(a)) := a$ and hence we get a coefficient map c_a that is independent of the other values of s . For example, we can take $a := [\gamma] \in \pm \pi_1(M) / \sim$ with $\gamma \in \pi_1(M)$ to get c_γ .

Definition 2.17 For $\gamma \in \pi_1(M)$, write $\lambda(f, g)_\gamma := c_\gamma(\lambda(f, g))$. This quantity lies in \mathbb{Z} (respectively $\mathbb{Z}/2$) when $|\gamma|$ lies in B_2 (respectively B_1), or equivalently when $[\gamma]$ has infinite order (respectively order 2) in $\Gamma_{f,g}$. The values do not depend on the choice of s and satisfy $c_{\gamma_1} = -c_{\gamma_2}$ whenever $[\gamma_1] = -[\gamma_2]$.

The following can be proven using Proposition 2.10 (see eg [Freedman and Quinn 1990, Section 1.7; Powell and Ray 2021b] for the case of discs and spheres).

Proposition 2.18 Let $f: \Sigma \rightarrow M$ and $g: \Sigma' \rightarrow M$ be based maps that are transverse to one another. The intersection number $\lambda(f, g)$ is preserved by homotopies that are ambient isotopies near $\partial \Sigma \sqcup \partial \Sigma'$.

Remark 2.19 The geometric definition of λ given above has a well-known algebraic version in the case that f and g correspond to classes in $H_2(M, \partial M; \mathbb{Z}[\pi_1(M)])$. This extends to the case of positive genus, as we now sketch. We restrict ourselves to the case that M , Σ and Σ' are closed and oriented for convenience.

Choose a basepoint in the universal cover of M , lifting the basepoint of M . The maps f and g lift uniquely (with respect to this choice of basepoint) to covers \widehat{M} and \widehat{M}' , corresponding to the subgroups $f \bullet (\pi_1(\Sigma))$ and $g \bullet (\pi_1(\Sigma'))$, respectively. These lifts represent classes

$$[f] \in H_2(\widehat{M}; \mathbb{Z}) \cong H_2(M; \mathbb{Z}[\pi_1(M)/f \bullet \pi_1(\Sigma)])$$

and

$$[g] \in H_2(\widehat{M}'; \mathbb{Z}) \cong H_2(M; \mathbb{Z}[\pi_1(M)/g \bullet \pi_1(\Sigma')]).$$

Then we have

$$\text{PD}^{-1}([f]) \smile \text{PD}^{-1}([g]) \in H^4(M; \mathbb{Z}[\pi_1(M)/f \bullet \pi_1(\Sigma)] \otimes_{\mathbb{Z}} \mathbb{Z}[\pi_1(M)/g \bullet \pi_1(\Sigma')]).$$

By Poincaré duality, this yields an element in

$$H_0(M; \mathbb{Z}[\pi_1(M)/f_\bullet\pi_1(\Sigma)] \otimes_{\mathbb{Z}} \mathbb{Z}[\pi_1(M)/g_\bullet\pi_1(\Sigma')]),$$

which is isomorphic as an abelian group to

$$\mathbb{Z}[f_\bullet\pi_1(\Sigma)\backslash\pi_1(M)] \otimes_{\mathbb{Z}[\pi_1(M)]} \mathbb{Z}[\pi_1(M)/g_\bullet\pi_1(\Sigma')].$$

Here $\mathbb{Z}[f_\bullet(\pi_1(\Sigma))\backslash\pi_1(M)]$ denotes $\mathbb{Z}[\pi_1(M)/f_\bullet\pi_1(\Sigma)]$ considered as a right $\mathbb{Z}[\pi_1(M)]$ -module.

Finally, we have the isomorphism

$$\begin{aligned} \mathbb{Z}[f_\bullet\pi_1(\Sigma)\backslash\pi_1(M)] \otimes_{\mathbb{Z}[\pi_1(M)]} \mathbb{Z}[\pi_1(M)/g_\bullet\pi_1(\Sigma')] &\rightarrow \mathbb{Z}[f_\bullet\pi_1(\Sigma)\backslash\pi_1(M)/g_\bullet\pi_1(\Sigma')], \\ [a] \otimes [b] &\mapsto [ab] \quad \text{for } a, b \in \pi_1(M). \end{aligned}$$

We shall not prove that this formulation agrees with the geometric definition.

2.3 Self-intersection numbers

Next we turn to the *self-intersection number* for a based generic immersion $f: \Sigma \looparrowright M$ of a connected surface Σ , with whisker v_f . The definition of μ , given below, is similar to that of λ in the previous subsection, except that there is no longer a clear choice of which sheet to consider first at a given double point. Consequently, the values of μ lie in a further quotient of the group $\Gamma_{f,f}$ from Definition 2.12.

We write $f \pitchfork f \subseteq M$ for the set of double points of f . We record the self-intersections of f by the sum of signed group elements

$$\mu(f) := \sum_{p \in f \pitchfork f} \varepsilon(p) \cdot \eta(p)$$

as follows:

- For $p = f(x_1) = f(x_2)$ for $x_1 \neq x_2 \in \Sigma$, let γ_1^p and γ_2^p be paths in Σ from the basepoint to x_1 and x_2 , respectively.
- The sign $\varepsilon(p) \in \{\pm 1\}$ is defined as follows. Transport the local orientation of Σ at the basepoint to x_1 along γ_1^p , and along γ_2^p to x_2 . This induces a local orientation at p . Another local orientation is obtained by transporting the local orientation at the basepoint of M to p along the concatenated path $v_f * (f \circ \gamma_2^p)$. We define $\varepsilon(p) = 1$ when the two local orientations match at p , and -1 otherwise.
- The element $\eta(p) \in \pi_1(M)$ is given by the concatenation $v_f * (f \circ \gamma_1^p) * (f \circ \gamma_2^p)^{-1} * v_f^{-1}$.

There is a similar discussion about homotopy invariance of $\mu(f)$ as for $\lambda(f, g)$ earlier: homotopies f_t that are generic immersions for all t preserve the formal sum of signed elements but finger moves and Whitney moves (of f with itself) create pairs $-\eta(p) + \eta(p)$, so it is convenient to use abelian groups. This takes care of regular homotopies of f but there is an additional subtlety for cusp homotopies f_t , where there is exactly one time t_0 for which f_{t_0} is not an immersion. These issues will be discussed carefully below.

For simply connected Σ , the self-intersection invariant $\mu(f)$ is well defined in the quotient (as an abelian group) of $\mathbb{Z}[\pi_1(M)]$ obtained by introducing the relators

$$\gamma - w^M(\gamma) \cdot \gamma^{-1}$$

for all $\gamma \in \pi_1(M)$. For general Σ , the quantity $\mu(f)$ is well defined in the abelian group

$$(2-3) \quad \Gamma_f := \Gamma_{f,f} / \langle \gamma - w^M(\gamma) \cdot \gamma^{-1} \rangle,$$

ie in this quotient of $\Gamma_{f,f}$ from Definition 2.12. Here, as above, $w^M : \pi_1(M) \rightarrow \{\pm 1\}$ is the orientation character.

We now change our notation slightly from the discussion of $\lambda(f, g)$ in order to work in this further quotient. Let \sim denote the equivalence relation on $\pm\pi_1(M)$ induced by the composition

$$\pm\pi_1(M) \hookrightarrow \mathbb{Z}[\pi_1(M)] \twoheadrightarrow \Gamma_f$$

sending $a \mapsto [a]$ and let $|\pi_1(M)|$ be the quotient of $\pm\pi_1(M)$ obtained by identifying γ_1 and γ_2 whenever $[\gamma_1] = \pm[\gamma_2] \in \Gamma_f$. Then we obtain the following analogues of Lemmas 2.14 and 2.16.

Lemma 2.20 *Let $\gamma_1, \gamma_2 \in \pi_1(M)$. One of the relations $[\gamma_1] = \pm[\gamma_2]$ holds if and only if γ_1 and γ_2 represent the same element in the quotient of the double coset by inversion. In other words, the identity map induces a bijection*

$$|\pi_1(M)| \leftrightarrow (f \bullet \pi_1(\Sigma) \backslash \pi_1(M) / f \bullet \pi_1(\Sigma)) / \approx,$$

where \approx is the equivalence relation identifying γ and γ^{-1} for all $\gamma \in \pi_1(M)$.

We write $|\gamma| := \text{pm}(\gamma)$ for the quotient map $\text{pm} : \pm\pi_1(M) / \sim \rightarrow |\pi_1(M)|$. Again, pm has fibres of order 1 or 2 and we decompose $|\pi_1(M)|$ as a disjoint union $B_1 \sqcup B_2$ according to this distinction as before. Choose a section $s : |\pi_1(M)| \rightarrow \pm\pi_1(M) / \sim$ of pm . As before, for $b \in B_2$ we denote the elements of the fibre by $\text{pm}^{-1}(b) = \{s(b), -s(b)\}$.

Lemma 2.21 *The abelian group Γ_f is a direct sum $\Gamma_f = FA \oplus V$ of a free abelian group FA on the set $s(B_2) \subseteq \pm\pi_1(M) / \sim \subseteq \Gamma_f$ and a $\mathbb{Z}/2$ -vector space V with basis $s(B_1) = \text{pm}^{-1}(B_1) \subseteq \pm\pi_1(M) / \sim$.*

Reading off the coefficients in this decomposition gives homomorphisms $c_s(b) : \Gamma_f \rightarrow \mathbb{Z}$ for $b \in B_2$, and $c_b : \Gamma_f \rightarrow \mathbb{Z}/2$ for $b \in B_1$, leading to a decomposition of Γ_f as a direct sum of copies of \mathbb{Z} and $\mathbb{Z}/2$.

Proof The proof is analogous to that of Lemma 2.16. □

Remark 2.22 If M and Σ are oriented, then

$$\Gamma_f \cong \mathbb{Z}[f \bullet \pi_1(\Sigma) \backslash \pi_1(M) / f \bullet \pi_1(\Sigma)] / \langle \gamma - \gamma^{-1} \rangle = \mathbb{Z}[(f \bullet \pi_1(\Sigma) \backslash \pi_1(M) / f \bullet \pi_1(\Sigma)) / \approx]$$

is free abelian. In this case, $B_1 = \emptyset$ and $B_2 = |\pi_1(M)| = (f \bullet \pi_1(\Sigma) \backslash \pi_1(M) / f \bullet \pi_1(\Sigma)) / \approx$ by Lemma 2.20.

Consider instead $f : \mathbb{R}P^2 \looparrowright M$ with $f_\bullet(\mathbb{R}P^1) = 1$, without any assumption on M . Then

$$\Gamma_f \cong (\mathbb{Z}/2)[\pi_1(M)] / \langle \gamma - w^M(\gamma) \cdot \gamma^{-1} \rangle = (\mathbb{Z}/2)[\pi_1(M) / \approx].$$

As in Remark 2.13, this is 2-torsion due to the relations $\gamma = w^{\mathbb{R}P^2}(\mathbb{R}P^1) \cdot f_\bullet(\mathbb{R}P^1) * \gamma = -\gamma$. In this case, $B_1 = |\pi_1(M)| = \pi_1(M) / \approx$ and $B_2 = \emptyset$.

As before, the subgroups FA and V of Γ_f do not depend on the choice of s ; only the basis of FA does. As a consequence, the coefficient maps $c_{s(b)}$ only depend on s up to sign and satisfy $c_{-s(b)} = -c_{s(b)}$. Given $a \in \pm\pi_1(M) / \sim$, we may again take $s(\text{pm}(a)) := a$ to get c_a and, in particular, c_γ for $\gamma \in \pi_1(M)$ independent of the choice of s at other points. This gives the following definition.

Definition 2.23 For $\gamma \in \pi_1(M)$, we write $\mu(f)_\gamma := c_\gamma(\mu(f))$. This quantity lies in \mathbb{Z} (respectively $\mathbb{Z}/2$) when $|\gamma|$ lies in B_2 (respectively B_1), or equivalently when $[\gamma]$ has infinite order (respectively order 2) in Γ_f . The values do not depend on the choice of s and satisfy $c_{\gamma_1} = -c_{\gamma_2}$ whenever $[\gamma_1] = -[\gamma_2]$.

We focus on $\mu(f)_1$, which plays an important role in the distinction between the homotopy class and regular homotopy class of f , as we will discuss in the next subsection. In the usual case, where Σ is simply connected, $\mu(f)_1 \in \mathbb{Z}$. However, in general, $\mu(f)_1$ may lie in either \mathbb{Z} or $\mathbb{Z}/2$. The following lemma gives the precise conditions determining the home of $\mu(f)_1$.

Lemma 2.24 Let $f : \Sigma \looparrowright M$ be a based, generic immersion, with whisker v . Recall that the map $f_\bullet : \pi_1(\Sigma) \rightarrow \pi_1(M)$ is given by $\alpha \mapsto v * (f \circ \alpha) * v^{-1}$.

If w^Σ is trivial on $\ker(f_\bullet)$ and w^M is trivial on $\text{Im}(f_\bullet)$, then $[1] \in \Gamma_f$ has infinite order and thus $\mu(f)_1 \in \mathbb{Z}$. Otherwise, $[1]$ has order 2 and $\mu(f)_1 \in \mathbb{Z}/2$.

Proof By definition, for $1 \in \pi_1(M)$, we know that $[1] \in \Gamma_{f,f}$ has order 2 precisely if

- (i) there exists $\alpha, \beta \in \pi_1(\Sigma)$ such that $f_\bullet(\alpha) * f_\bullet(\beta) = f_\bullet(\alpha * \beta) = 1$ and $w^\Sigma(\alpha)w^\Sigma(\beta)w^M(f_\bullet(\beta)) = w^\Sigma(\alpha * \beta)w^M(f_\bullet(\beta)) = -1$, or
- (ii) there exists $\delta \sim 1$ where δ has order two in $\pi_1(M)$ and $w^M(\delta) = -1$.

Suppose that w^Σ is trivial on $\ker(f_\bullet)$ and w^M is trivial on $\text{Im}(f_\bullet)$. Then the first case (i) cannot happen since, if $f_\bullet(\alpha * \beta) = 1$, then $\alpha * \beta \in \ker(f_\bullet)$ so $w^\Sigma(\alpha * \beta)w^M(f_\bullet(\beta)) = 1 \cdot 1 = 1$. Similarly, (ii) cannot happen: suppose δ has order two in $\pi_1(M)$ and $\delta \sim 1$. Then, by definition, $\delta = f_\bullet(\alpha) * 1 * f_\bullet(\beta) = f_\bullet(\alpha * \beta)$ in $\pi_1(M)$, for some $\alpha, \beta \in \pi_1(M)$. In particular, $\delta \in \text{Im}(f_\bullet)$, and so again $w^M(\delta) = 1$ by hypothesis, contradicting (ii). Therefore, $[1]$ has infinite order, as claimed.

Now suppose there is some $\alpha \in \pi_1(\Sigma)$ with $w^M(f_\bullet(\alpha)) = -1$. Then we have $f_\bullet(\alpha^{-1}) * 1 * f_\bullet(\alpha) = 1$ and $w^\Sigma(\alpha^{-1})w^\Sigma(\alpha)w^M(f_\bullet(\alpha)) = -1$, so $[1] \sim -[1]$ and $[1]$ has order two.

Finally, suppose that there is some $\alpha \in \ker(f_\bullet)$ with $w^\Sigma(\alpha) = -1$. Then we have $f_\bullet(\alpha) * 1 * f_\bullet(1_\Sigma) = 1$ and $w^\Sigma(\alpha)w^\Sigma(1_\Sigma)w^M(f_\bullet(1_\Sigma)) = -1$, where 1_Σ denotes the trivial element in $\pi_1(\Sigma)$. Then, again, we have $[1] \sim -[1]$ and $[1]$ has order two. □

As with Proposition 2.18, the proof of the following proposition is virtually identical to the case of discs and spheres using Proposition 2.10 (see eg [Powell and Ray 2021b]), and we leave it for the interested reader.

Proposition 2.25 *Let $f: \Sigma \looparrowright M$ be a based generic immersion. The self-intersection number $\mu(f)$ is preserved under regular homotopies that are ambient isotopies near $\partial\Sigma$.*

In this and the previous subsection, we have considered intersection and self-intersection numbers of connected surfaces. By combining these invariants, we can define the conglomerate notion of self-intersection number for disconnected surfaces $F = \{f_i\}_{i=1}^m: \Sigma \rightarrow M$, as considered in Definition 1.3:

$$\mu(F) := \sum_{i < j} \lambda(f_i, f_j) + \sum_i \mu(f_i) \in \bigoplus_{i < j} \Gamma_{f_i, f_j} \oplus \bigoplus_i \Gamma_{f_i}.$$

Propositions 2.18 and 2.25 imply that $\mu(F)$ is preserved under regular homotopies of F that are ambient isotopies near $\partial\Sigma$.

2.4 Whitney discs

A Whitney move cancels a pair of double points of a generic immersion $F: \Sigma \looparrowright M$ as in Convention 1.1, provided all the assumptions on the guiding Whitney disc are satisfied. In our setting, where Σ and M need be neither simply connected nor orientable, this requires some care. We start with the notion of arcs A and A' pairing double points p and q , and the corresponding notion of (p, q, A, A') having *opposite sign*.

Definition 2.26 Let $f: \Sigma \rightarrow M$ and $g: \Sigma' \rightarrow M$ be based maps that either intersect transversely, or $f = g$ and f is a generic immersion. We say that two points $p, q \in f \pitchfork g \subseteq M$ are *paired by arcs* if we equip them with the extra data of an arc $A: [0, 1] \rightarrow \Sigma$ from $f^{-1}(p)$ to $f^{-1}(q)$ and an arc A' in Σ' from $g^{-1}(q)$ to $g^{-1}(p)$. In the case that $f = g$, we require that each point in $f^{-1}(p)$ and in $f^{-1}(q)$ is the endpoint of precisely one of the arcs A and A' , ie A and A' lie in distinct sheets at both p and q .

With the extra data of the arcs A and A' , we can make sense of whether two intersection points that are paired by arcs have opposite sign.

Definition 2.27 Let $f: \Sigma \rightarrow M$ and $g: \Sigma' \rightarrow M$ be based maps that either intersect transversely, or $f = g$ and f is a generic immersion. Two intersection points $p, q \in f \pitchfork g \subseteq M$ paired by arcs A in Σ and A' in Σ' have *opposite sign* if the following holds. Fix local orientations of Σ at $f^{-1}(p)$ and of Σ' at $g^{-1}(p)$. This choice induces a local orientation of M at p . Transport the local orientation of Σ from $f^{-1}(p)$ to $f^{-1}(q)$ along A , and the local orientation of Σ' from $g^{-1}(p)$ to $g^{-1}(q)$ along A' . This gives a local orientation of M at the point q . Compare this with the local orientation on M at q induced by transporting the local orientation from p to q along the arc $f \circ A$. If these orientations disagree, then the points p and q are said to have opposite sign (with respect to A and A'), and otherwise they are said to have the same sign. The dependence on the choice of arcs A and A' is sometimes neglected.

Note that double points having the same sign could be “paired” by an embedded disc, but this does not mean that a Whitney move using this disc is possible, because the required section of the normal bundle of the disc is not available; in this case, any rank one subbundle of the normal bundle of the disc, restricted to the boundary, that is tangent to one sheet of Σ and normal to the other sheet turns out to be a Möbius bundle. So one does not study such discs and assumes that a Whitney disc always pairs two double points of opposite sign.

In the setting of based transverse maps $f \neq g$, with Σ and Σ' connected, recall from Section 2.2 that $\lambda(f, g)$ is a sum of terms $\varepsilon(p) \cdot \eta(p)$, one for each double point $p \in f \pitchfork g$, with $\eta(p) \in \pi_1(M)$ and $\varepsilon(p) \in \{\pm 1\}$. This sum is well defined in the abelian group $\Gamma_{f,g}$ and each signed group element $a \in \pm\pi_1(M)$ represents a unique element $[a] \in \Gamma_{f,g}$. The same proof as in the case of simply connected surfaces [Powell and Ray 2021b, Proposition 11.10] yields the following result.

Lemma 2.28 *Let Σ and Σ' be compact connected surfaces and let $f : \Sigma \rightarrow M$ and $g : \Sigma' \rightarrow M$ be based maps with transverse double points $p, q \in f \pitchfork g \subseteq M$. Then $[\varepsilon(p) \cdot \eta(p) + \varepsilon(q) \cdot \eta(q)] = 0 \in \Gamma_{f,g}$ if and only if p and q can be paired by arcs $A \subseteq \Sigma$ and $A' \subseteq \Sigma'$ such that*

- (i) *the closed loop $f \circ A \cup_{p,q} g \circ A'$ is null-homotopic in M , and*
- (ii) *the points p and q have opposite sign with respect to the arcs A and A' .*

If (i) and (ii) are satisfied for p and q , we say that $W : D^2 \rightarrow M$ is a (map of a) *Whitney disc pairing p and q* if its boundary is the closed loop in (i), the union of its two *Whitney arcs* $f \circ A$ and $g \circ A'$. We leave it to the reader to formulate the analogous notion for a pair of transverse self-intersection points of $f : \Sigma \rightarrow M$. This gives rise to the following corollary to Lemma 2.28.

Corollary 2.29 *Let Σ and Σ' be compact connected surfaces and let $f : \Sigma \rightarrow M$ and $g : \Sigma' \rightarrow M$ be transverse based maps. Then $\lambda(f, g) = 0$ if and only if all intersection points between f and g can be paired by maps of Whitney discs.*

Moreover, a based generic immersion $f : \Sigma \looparrowright M$ satisfies $\mu(f) = 0$ if and only if all self-intersection points of f can be paired by maps of Whitney discs.

Note that, by the geometric Casson lemma (Lemma 4.2), the vanishing of $\lambda(f, g)$ is equivalent to the existence of a regular homotopy of f and g that makes their images disjoint, at the cost of introducing self-intersections in f and g . There is no analogue of this argument if $\mu(f) = 0$ by the failure of the Whitney trick in dimension 4, as for example exhibited by the secondary embedding obstruction $\text{km}(f)$ (see Section 3).

By the geometric characterisation in Corollary 2.29, it is meaningful to refer to f and g having trivial intersection number, and to f having trivial self-intersection number, without using a basing.

The analogue of the characterisation in the second part of Corollary 2.29 holds for generic immersions $F = \{f_i\}_{i=1}^m : \Sigma \looparrowright M$ from Convention 1.1, ie for compact but possibly disconnected domains, if we use Definition 1.3 from the introduction for the self-intersection number $\mu(F)$.

Corollary 2.30 *Let $F : (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. Then $\mu(F) = 0$ if and only if the double points of F can be paired by maps of Whitney discs.*

Proof This is a direct consequence of Propositions 2.18 and 2.25 and Corollary 2.29 because every double point of F is either a self-intersection point of a component f_i or an intersection point between distinct components f_i and f_j (where we can assume that $i < j$). Note that both cases represent self-intersection points of F . \square

Collections of Whitney discs as above may be assumed to be *convenient* in the following sense (see eg [Freedman and Quinn 1990, Section 1.4; Powell and Ray 2021b]).

Definition 2.31 *Let $F : (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. A *convenient* collection of Whitney discs for F is a collection of framed, generically immersed Whitney discs pairing all the double points of F , with interiors transverse to F and with disjointly embedded boundaries. A collection of arcs in $F(\Sigma)$ is called a collection of *Whitney arcs* if the union of the arcs is the boundary of a convenient collection of Whitney discs.*

By pushing double points of a convenient collection across the boundaries of Whitney discs [Powell and Ray 2021b, Figure 11.4], we may further assume that all Whitney discs are pairwise disjoint and embedded. However, the (resulting and preexisting) intersections between the original surface F and the Whitney discs can in general not be removed, as detected by the secondary invariant $\text{km}(F)$.

2.5 Homotopy versus regular homotopy of generic immersions

Let $f : \Sigma \looparrowright M$ be a generic immersion. Local orientations of M and Σ determine a local orientation of νf . Hence, given a framing of $f|_{\partial\Sigma}$, one can define a relative Euler class of the normal bundle νf in $H^2(\Sigma, \partial\Sigma; \mathbb{Z}^{w_1(\nu f)})$. If $f^*(w_1(M)) = w_1(\nu f) + w_1(T\Sigma) = 0$ then the local orientation of Σ determines a Poincaré duality isomorphism from this twisted cohomology group to \mathbb{Z} , and we denote the resulting integer by $e(\nu f)$. Note that $e(\nu f)$ does not depend on the local orientation of Σ but only on the local orientation of M . If $f^*(w_1(M)) \neq 0$ then there is still a mod 2 normal Euler number, which we also denote by $e(\nu f) \in \mathbb{Z}/2$.

A useful interpretation of $e(\nu f)$ is as follows. A vector in \mathbb{R}^2 together with the framing of $f|_{\partial\Sigma}$ determines a nonvanishing section of νf on $f(\partial\Sigma)$. Extend this to a section of νf over all of $f(\Sigma)$, transverse to the zero section. Then $e(\nu f)$ counts, with sign, the number of zeros of the section, in \mathbb{Z} or $\mathbb{Z}/2$ as appropriate.

Next we give an extension of [Powell et al. 2020, Theorem 1.2] from the simply connected to the general setting, restricting ourselves to the case of connected Σ for convenience. We note that [Powell et al. 2020, Theorem 1.2] was based on [Freedman and Quinn 1990, Lemma 1.2 and Proposition 1.6], but that the latter proposition was not proven in [Freedman and Quinn 1990].

By the following theorem, in some cases, for example when M is orientable, $e(vf) \in \mathbb{Z}$ is an additional invariant of regular homotopy classes of immersions. It changes by ± 2 during a cusp homotopy (see eg [Conant et al. 2012b, Figure 19]) and hence there can be infinitely many regular homotopy classes of immersions, all of which are homotopic as continuous maps.

In Theorem 2.32, in the case that Σ has nontrivial boundary, we fix a framing on the embedding $f|_{\partial\Sigma}$, in order to define the relative Euler number $e(v\tilde{f})$ for \tilde{f} any generic immersion homotopic to f .

Theorem 2.32 *Let Σ be a compact, connected surface and let M be a 4-manifold. Then the inclusion of the subspace of generic immersions $\text{Imm}(\Sigma, M)$ in the space of all continuous maps induces a map*

$$\frac{\text{Imm}(\Sigma, M)}{\{\text{regular homotopy}\}} \xrightarrow{i} [\Sigma, M]_{\partial},$$

where $[\Sigma, M]_{\partial}$ denotes the set of homotopy classes of continuous maps that restrict on $\partial\Sigma$ to embeddings disjoint from the image of the interior of Σ .

- (1) i is surjective.
- (2) The fibres of i are related by cusp homotopies. More precisely, suppose that f and g are homotopic generic immersions. Then we can add cusps to f and g , to obtain f' and g' , respectively, such that f' and g' are regularly homotopic.
- (3) For every $f \in [\Sigma, M]_{\partial}$, there is a bijection

$$i^{-1}(f) \cong \begin{cases} 2\mathbb{Z} & \text{if } f^*(w_1(M)) = 0 \text{ and } w_2(v\tilde{f}) = 0, \\ 2\mathbb{Z} + 1 & \text{if } f^*(w_1(M)) = 0 \text{ and } w_2(v\tilde{f}) = 1, \\ \mathbb{Z}/2 & \text{otherwise,} \end{cases}$$

where $v\tilde{f}$ is a normal bundle for \tilde{f} , a generic immersion in $i^{-1}(f)$. When $f^*(w_1(M)) = 0$, the bijection is given by

$$\tilde{f} \mapsto e(v\tilde{f}).$$

Otherwise, the bijection is given by

$$\tilde{f} \mapsto \mu(\tilde{f})_1 \in \mathbb{Z}/2.$$

- (4) If $f^*(w_1(M)) = 0$ and $w_1(\Sigma)|_{\ker(f_{\bullet})} = 0$ for \tilde{f} a generic immersion in $i^{-1}(f)$, the quantities $\mu(\tilde{f})_1$ and $e(v\tilde{f})$ are related by the formula

$$\lambda(\tilde{f}, \tilde{f})_1 = 2\mu(\tilde{f})_1 + e(v\tilde{f}) \in \mathbb{Z}$$

and so $\mu(\tilde{f})_1 \in \mathbb{Z}$ also detects the regular homotopy class of $\tilde{f} \in i^{-1}(f)$.

While we prefer the upcoming direct argument analysing singularities, Theorem 2.32 could in principle also be proven via Smale–Hirsch immersion theory, which has a version in the topological category. The main novelty of the theorem is that we give precise conditions in terms of the Stiefel–Whitney classes to control how large the fibres of i are, and which invariants detect them.

Proof By Proposition 2.10(1), the map is surjective. That is, every homotopy class contains a generic immersion. This proves (1).

For (2), note that, if f and g are homotopic generic immersions, then, by Proposition 2.10(2), there exists a generic homotopy H between them, which by definition is a sequence of ambient isotopies, finger moves, Whitney moves and cusp homotopies. We can modify H so that there are real numbers $t_1 < t_2 \in [0, 1]$ such that the singularities of H in $[0, t_1]$ only consist of cusp homotopies that create double points, the singularities in $[t_1, t_2]$ only consist of finger moves and Whitney moves, and those in $[t_2, 1]$ only consist of cusp homotopies that remove double points. The statement then follows by taking $f' := H_{t_1}$ and $g' := H_{t_2}$.

To achieve this modification, note that we can bring all the creating cusp singularities forward, so that they occur earlier, and we can delay all the removing cusps. To arrange for a creating cusp to be rearranged earlier than a finger or Whitney move, choose an arc in the image of H starting from $H_t(\Sigma)$ for some $t \in (0, t_1)$, and ending at the cusp, which intersects each level in a point and is disjoint from all Whitney arcs and double points. The homotopy can then be altered in a neighbourhood of this arc so that the cusp singularity occurs at time t . Delaying a removing cusp is the same procedure but with the direction of time reversed. This completes the proof of (2).

The proof of (3) splits naturally into two cases.

Case 1 $f^*(w_1(M)) = 0$.

As noted in Section 2, the sign of an intersection point is not always well defined. Nevertheless, in the case that $f^*(w_1(M)) = 0$, the sign of a cusp homotopy is well defined. The key point is that a cusp not only specifies a double point p but also an arc between the preimages of p . In the case that $f^*(w_1(M)) = 0$, using this path, the sign of the double point p is well defined, independent of the choice of path transporting the local orientation at the basepoint to the double point. Thus in this setting we define the *sign* of a cusp to be the sign of the double point it creates or removes. We will use the terminology of *creating cusps* for cusps that create a double point and *removing cusps* for those that remove a double point.

Since $f^*(w_1(M)) = 0$, $e(v\tilde{f})$ is defined in \mathbb{Z} for any generic immersion \tilde{f} homotopic to f . Recall that $w_2(v\tilde{f}) \equiv e(v\tilde{f}) \pmod{2}$. Since regularly homotopic generic immersions have equal Euler numbers, the map in the theorem statement is well defined on equivalence classes in the domain of i . Note that a cusp homotopy changes $e(v\tilde{f})$ by 2 or -2 , depending on the sign of the cusp and whether it is a creating or a removing cusp. So every element of $2\mathbb{Z}$ or $2\mathbb{Z} + 1$, depending on $w_2(v\tilde{f})$, can be realised as the Euler number of a generic immersion in $i^{-1}(f)$.




Figure 3: A schematic picture showing how a removing cusp singularity and creating cusp singularity with the same sign can be cancelled. In each of (a), (b) and (c), a homotopy is traced out in the direction of t . At every time t , except the times of the cusp singularities, we depict an arc of a generic immersion homotopic to f . (a) Two cusp singularities are shown: a removing cusp occurring first, followed by a creating cusp of the same sign. (b) Modify the homotopy, delaying the removing cusp until it coincides with the creating cusp. This involves choosing an arc in Σ joining the two cusp points. (c) A further local modification removes the two cusps.

To complete the proof when $f^*(w_1(M)) = 0$, it remains to show injectivity. We will show that, given a generic homotopy between generic immersions with equal Euler numbers, we can modify the homotopy to cancel cusps, until we are left with a regular homotopy.

First note that, when we have a removing cusp, and later in the homotopy we have a creating cusp with the same sign, we can cancel these two cusps along a level-preserving path in the homotopy, as indicated in Figure 3.

However, this is not sufficient. We also have to show that we can also cancel cusps given

- (i) two creating cusps of opposite sign, or two removing cusps of opposite sign; or
- (ii) a creating cusp paired with a later removing cusp, both of the same sign.

Suppose that we have a generic homotopy H between generic immersions with equal Euler numbers consisting of two creating cusp homotopies of opposite sign, as in (i). Create a self-homotopy H_0 of the starting immersion, ie the immersion at $t = 0$, consisting of a trivial finger move together with two removing cusps for the double points created by the finger move, as shown in Figure 4. Then concatenate H_0 with the original homotopy H . The new homotopy can be modified as in Figure 3 to cancel the removing cusps in H_0 and the creating cusps in H , leaving only the finger move behind. An analogous argument shows how to cancel two removing cusps of opposite sign, this time concatenating at the end of H .

Similarly, for the situation in (ii), suppose that we have a generic homotopy H between generic immersions with equal Euler numbers consisting of a creating cusp and a later removing cusp of the same sign. Again




Figure 4: A schematic picture showing a self-homotopy consisting of a trivial self-finger move followed by two removing cusps. The homotopy is traced out in the direction of t . At every time t , except at the times of the cusp and finger move singularities, we depict an arc of a generic immersion homotopic to f . In red we show the arc of self-intersections of f ; note that it starts at one cusp singularity and ends at the other.

we construct the self-homotopy H_0 and concatenate with H . In the result, we use the procedure from Figure 3 to cancel the creating cusp in H with one of the removing cusps in H_0 . This entire operation has so far replaced a cusp with a cusp of opposite sign and direction. As before we can repeat the operation at the end of the homotopy to replace the removing cusp with a creating one, also with the opposite sign. Thus when we have a creating cusp with a later removing cusp of the same sign, we can replace both by cusps of opposite sign and direction. Since now the removing cusp happens before the creating cusp, the two can be cancelled and we are done with case (ii). This completes the proof of (3) in the case that $f^*(w_1(M)) = 0$.

Case 2 $f^*(w_1(M)) \neq 0$.

Note that a cusp homotopy changes $\mu(\tilde{f})_1 \in \mathbb{Z}/2$ by one. So both values of $\mathbb{Z}/2$ can be realised within the homotopy class. To show injectivity in this case, we have to show that we can cancel cusps in a homotopy in arbitrary pairs. First use the trading argument above to get all the removing cusps before the creating cusps in the homotopy. Then, for any pair of cusps, one removing and one creating, choose some level-preserving path in the homotopy between the first and the second cusp, and restrict to a small disc containing the path. If they have the same sign with respect to this disc, cancel the two cusps as before.

If they have opposite signs, change the choice of the arc to arrange that the union of the new arc and the old arc maps nontrivially under $w_1(M)$. Such an arc exists since $f^*(w_1(M))$ is nontrivial and Σ is connected. With this new choice, the signs of the cusps in the disc become the same and we can again cancel the cusps. This completes the proof of both halves of (3).

Finally, for (4), note that, if $f^*(w_1(M)) = 0$ and $w_1(\Sigma)|_{\ker(f_\bullet)} = 0$, then, by Lemma 2.24, $\mu(\tilde{f})_1$ is well defined in \mathbb{Z} . By the discussion above the statement of the theorem, $e(v\tilde{f})$ is also well defined in \mathbb{Z} . In this case, the formula

$$\lambda(\tilde{f}, \tilde{f})_1 = 2\mu(\tilde{f})_1 + e(v\tilde{f}) \in \mathbb{Z}$$

holds by the proof of the corresponding fact for discs and spheres (see eg [Powell and Ray 2021b, Proposition 11.8]). Any cusp homotopy leaves $\lambda(\tilde{f}, \tilde{f})_1$ unchanged, while it changes $\mu(\tilde{f})_1$ by ± 1 . By

the formula, it changes $e(v\tilde{f})$ by ∓ 2 . Thus, if \tilde{f} and \tilde{f}' are generic immersions homotopic to f , then $e(v\tilde{f}) = e(v\tilde{f}')$ if and only if $\mu(\tilde{f})_1 = \mu(\tilde{f}')_1$. Hence, (4) follows from (3). \square

3 Secondary embedding obstructions

The Whitney trick implies that every map $F: S^n \rightarrow M^{2n}$ is homotopic to an embedding whenever M is a simply connected $2n$ -dimensional manifold and $n > 2$. In order to prove the failure of the Whitney trick in dimension 4, Kervaire and Milnor [1961] devised an obstruction that gave counterexamples to the above statement for $n = 2$. They showed that the homotopy class of $3 \cdot \mathbb{C}\mathbb{P}^1$ is not represented by an embedded sphere in $\mathbb{C}\mathbb{P}^2$. In a smooth, oriented, closed 4-manifold M , consider the formula

$$(3-1) \quad \theta(c) := \frac{1}{8}(c \cdot c - \sigma(M)) \pmod{2},$$

where the $\mathbb{Z}/2$ -reduction of $c \in H_2(M; \mathbb{Z})$ is Poincaré dual to $w_2(M)$ and $\sigma(M)$ is the signature of the intersection form $(x, y) \mapsto x \cdot y$ on $H_2(M; \mathbb{Z})$. In this setting, if c is represented by an embedded sphere, then $\theta(c) = 0$. Recall that, for a unimodular form ℓ and a characteristic element c , ie one satisfying $\ell(c, x) \equiv \ell(x, x) \pmod{2}$, the difference $\ell(c, c) - \sigma(\ell)$ is always divisible by 8. The condition on c being dual to $w_2(M)$ is stronger than being characteristic for the intersection form since the mod 2 intersection condition holds for all $x \in H_2(M; \mathbb{Z}/2)$, not just for integral homology classes. For example, if M is the Enriques surface (double covered by the $K3$ surface), then $\theta(0) \neq 0$, so 0 cannot be dual to $w_2(M)$, even though the intersection form on $H_2(M; \mathbb{Z})$ is even.

For the proof that θ is an embedding obstruction, Kervaire and Milnor added $1 - [F] \cdot [F]$ copies of $(\mathbb{C}\mathbb{P}^2, \mathbb{C}\mathbb{P}^1)$ to a proposed characteristic pair $(M, F: S^2 \hookrightarrow M)$, with F assumed to be an embedding, to obtain an embedded sphere with self-intersection number 1. Then they blow down that characteristic sphere to arrive at a spin manifold M' with $\sigma(M') = \sigma(M) + (1 - [F] \cdot [F]) - 1 = \sigma(M) - [F] \cdot [F]$. Rokhlin’s theorem [1952] — that the signature of a smooth, closed, spin 4-manifold is divisible by 16 — is equivalent to the original condition $\theta([F]) = 0$ in M .

The Kervaire–Milnor result also has consequences for spin manifolds, where it says that any (characteristic) homology class $c = 2b$ that is represented by an embedded sphere must satisfy $b \cdot b \equiv 0 \pmod{4}$. For example, $2\Delta \in \pi_2(S^2 \times S^2)$ is not represented by an embedding for Δ the diagonal 2-sphere.

For about a decade, it remained an open problem to find a combinatorial formula for $\theta(c)$ in terms of geometric representatives for c .

3.1 Combinatorial formulas: Rokhlin’s Arf invariant

Rokhlin [1972] picked an embedded representative $F: \Sigma \hookrightarrow M$ for $c \in H_2(M; \mathbb{Z})$ as above and assumed that $H_1(M; \mathbb{Z}/2)$ vanishes. Any simple closed curve r in (the image of) F then bounds a compact surface R in M . The reader should think of R as an “unoriented cap” and check that it has a relative Euler

number, just like a Whitney disc or an ordinary cap. Rokhlin then asserted that setting $q_F(r) := |\text{Int } R \pitchfork F|$ for R with vanishing relative Euler number defines a quadratic enhancement

$$q_F: H_1(\Sigma; \mathbb{Z}/2) \rightarrow \mathbb{Z}/2$$

that refines the mod 2 intersection form on Σ . Independence from the choice of R follows from F being dual to $w_2(M)$, in this setting using intersections of F with all classes of the form $[R \cup R'] \in H_2(M; \mathbb{Z}/2)$. Rokhlin stated that the Arf invariant $\text{Arf}(q_F)$ is equal to $\theta(c) = \theta([F])$. A nice consequence of this equality is that $\text{Arf}(q_F) = \theta(c)$ vanishes whenever c can be represented by an embedded sphere, because q_F is then defined on the zero vector space.

3.2 Combinatorial formulas: Freedman and Kirby's characteristic bordism

Using the same definitions, Freedman and Kirby [1978] proved Rokhlin's claims from above, on their way to a geometric proof of Rokhlin's original theorem. They worked with an arbitrary smooth, closed, oriented 4-manifold M , but before computing q_F they performed surgery on circles in M to arrange that $H_1(M; \mathbb{Z}/2) = 0$; alternatively, they could have made M simply connected and used discs for R , ie ordinary caps. They showed that $\text{Arf}(q_F)$ is invariant under "characteristic bordism", implying independence from the choice of surgeries, as well as establishing the equality $\text{Arf}(q_F) = \theta([F])$ by checking it on the generators of Ω_4^{char} . A different proof of $\text{Arf}(q_F) = \theta([F])$ was given in [Matsumoto 1986].

On a historical note, Freedman and Kirby wrote that they learnt these results from Casson and that they only heard of Rokhlin's results after finishing their paper. The Rokhlin method was extended to nonorientable characteristic surfaces in closed 4-manifolds in [Guillou and Marin 1980; Kirby and Taylor 2001].

3.3 Combinatorial formulas: Matsumoto's t -invariant

Matsumoto [1978], in the same proceedings as [Freedman and Kirby 1978], started with a spherical class $c \in \pi_2(M)$ and represented it by a generic immersion $F: S^2 \looparrowright M$ with $2g$ algebraically cancelling double points. He assumed that $H_1(M; \mathbb{Z}) = 0$, using this condition to find "Whitney surfaces", ie oriented surfaces R_1, \dots, R_g bounded by pairs of Whitney arcs in F . Again there is a relative Euler number and we may assume that every R_i has vanishing relative Euler number. Matsumoto proved that, if $[F] \in H_2(M; \mathbb{Z})$ is characteristic, then

$$(3-2) \quad \text{Arf}(q_F) = \sum_{i=1}^g |\text{Int } R_i \pitchfork F| =: t(F) \in \mathbb{Z}/2$$

by adding g tubes based at pairs of double points of F to turn it into an embedding of a surface Σ of genus g , where q_F is the quadratic enhancement defined above. The new surface has pairs of framed caps (D_i, R_i) , where D_i is a meridional disc of the i^{th} tube and hence has one interior intersection with

F , so $q_F(\partial D_i) = 1$. Since the boundaries of these caps form a hyperbolic basis of $H_1(\Sigma; \mathbb{Z}/2)$, the result follows from the usual formula

$$\text{Arf}(q_F) = \sum_{i=1}^g q_F(\partial D_i) \cdot q_F(\partial R_i) = \sum_{i=1}^g q_F(\partial R_i) = \sum_{i=1}^g |\text{Int } R_i \cap F|.$$

3.4 Summary of the secondary embedding obstructions from the 1970s

Given $[F] \in \pi_2(M)$ such that its Hurewicz image in $H_2(M; \mathbb{Z}/2)$ is Poincaré dual to $w_2(M)$, the above results show that

$$\theta([F]) = \text{Arf}(q_F) = t(F) \in \mathbb{Z}/2$$

is an obstruction to representing $[F]$ by an embedding $F: S^2 \hookrightarrow M$. Note that θ only depends on the homology class $h([F]) \in H_2(M; \mathbb{Z})$ by definition, whereas that is not clear for the other two invariants.

An attractive aspect of Matsumoto’s $t(F)$ is that it can be computed combinatorially from a generic immersion $F: S^2 \looparrowright M$. One argues directly that $t(F)$ is an obstruction to representing $[F] \in \pi_2(M)$ by an embedded sphere and independence of the choice of R_i comes from $[F]$ being characteristic.

Matsumoto’s formula was extended in a number of ways. For example, in recent work of Kasprowski, Land, Powell and Teichner [Kasprowski et al. 2017; 2021b; 2020] on the stable diffeomorphism classification of spin 4-manifolds, a version of Matsumoto’s t -invariant was used to compute the relevant Arf invariant. We describe further extensions presently.

It follows from topological transversality [Freedman and Quinn 1990, Section 9.5] that, in a smooth, closed, oriented 4-manifold M , the quantity $\theta(c)$ is also an obstruction to representing an element c as before by a topological — ie locally flat — embedding $F: S^2 \hookrightarrow M$. If M is not smooth, one adds the Kirby–Siebenmann invariant and then the formula

$$\theta_{\text{TOP}}(c) := \theta(c) + \text{ks}(M)$$

defines such an obstruction; see [Conant et al. 2012a, Introduction] for details. For example, it follows that the generator of $\pi_2(*\mathbb{C}\mathbb{P}^2)$ is not represented by an embedding. Historically speaking, these applications were not known at the time of publication of [Rokhlin 1972; Freedman and Kirby 1978; Matsumoto 1978].

In the following, we will return to considering topological manifolds and obstructions to topological embeddings.

3.5 Secondary obstructions to embedding genus zero surfaces with dual spheres

If Σ is a union of discs or spheres and $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ has algebraically dual spheres, then Freedman and Quinn [1990, Definition 10.8A] gave a version of Matsumoto’s t -invariant, calling it the Kervaire–Milnor invariant. Rather than restricting $H_1(M; \mathbb{Z})$ as in the discussions above, they assumed that $\mu(F) = 0$, ie that all double points of F can be paired by Whitney discs. They used the same

formula as in (3-2), but counted intersections with F° , restricting to the Whitney discs in a convenient collection \mathcal{W}° that pair double points of F° . They claimed that this mod 2 count, $t(F^\circ, \mathcal{W}^\circ)$ from Definition 1.5, is a secondary obstruction to representing F by an embedding. However, this is only true if F° is r -characteristic (Definition 5.5), as Stong's correction [1994] showed. Stong noticed that the choice of sheets for double points whose group elements have order 2 is related to immersed $\mathbb{R}P^2$ s in M . If F is dual to $w_2(M)$ then F is also r -characteristic, but not vice versa, so this obstruction is more generally defined than $\theta([F])$.

The embedding theorem for unions of discs and spheres [Freedman and Quinn 1990, Theorem 10.5], as corrected by Stong, says, in our notation, that, for good fundamental group $\pi_1(M)$, such an F is homotopic to a topological embedding if and only if there exists a convenient collection of Whitney discs \mathcal{W}° for the double points of F° such that $t(F^\circ, \mathcal{W}^\circ) = 0$. We give more details about the Freedman–Quinn–Stong embedding result in Section 5.

3.6 Secondary obstructions to embedding unions of spheres

Matsumoto's invariant $t(F)$ from (3-2) was extended to a secondary embedding obstruction in [Schneiderman and Teichner 2001] for $F = \{f_i\}_{i=1}^m$, not assuming dual spheres, where each $f_i: S^2 \looparrowright M$ is a generic immersion and assuming $\mu(F) = 0$ and that M is oriented. By counting interior intersections of F with a convenient collection \mathcal{W} of Whitney discs pairing the double points of F , and remembering group elements, signs and components of F , the authors defined an intersection count $\tau(F, \mathcal{W}) \in \mathcal{T}(\pi_1(M), m)$. Here $\mathcal{T}(\pi_1(M), m)$ is the abelian group given by the direct sum of $m + \binom{m}{2} + \binom{m}{3}$ copies of $\mathbb{Z}[\pi_1(M) \times \pi_1(M)]$. To obtain a secondary embedding obstruction, Schneiderman and Teichner [2001, Section 8] gave a list of relations such that the subgroup $\mathcal{R}(M, F) \leq \mathcal{T}(\pi_1(M), m)$ generated by these relations has the property that

$$\tau(f_1, \dots, f_m) = \tau(F) := [\tau(F, \mathcal{W})] \in \mathcal{T}(\pi_1(M), m) / \mathcal{R}(M, F)$$

does not depend on the choice of convenient collection \mathcal{W} . In our current language, the main result of that paper is that $\tau(F) = 0$ if and only if $\text{km}(F) = 0$ as in Definition 1.4. In the absence of dual spheres, $\tau(F) = 0$ does not imply that F is homotopic to an embedding. For example, there are obstructions from higher-order Whitney towers.

If F is r -characteristic then the augmentation map $\mathcal{E}: \mathcal{T}(\pi_1(M), m) \rightarrow \mathbb{Z}/2$, summing all possible coefficients, takes $\mathcal{R}(M, F)$ to zero and $\tau(F)$ to Matsumoto's $t(f_1 \# \dots \# f_m)$, for an arbitrary choice of interior connected sum of the $\{f_i\}_{i=1}^m$. Moreover, if F has algebraic dual spheres then \mathcal{E} induces an isomorphism of $\mathcal{T}(\pi_1(M), m) / \mathcal{R}(M, F)$ with either $\mathbb{Z}/2$ or 0, depending on whether F° is r -characteristic or not. This gives the relationship to Section 3.5.

3.7 Secondary embedding obstructions for arbitrary compact surfaces

It is likely possible to extend the invariant τ from Section 3.6 to arbitrary immersed compact surfaces, not just spheres. However, determining the analogue of $\mathcal{R}(M, F)$ would be a formidable task. In this

paper we take the first step, namely by defining the right notion of b -characteristic surfaces for which Matsumoto's invariant extends from spheres to a secondary embedding obstruction for arbitrary compact surfaces. We also generalise the work of Freedman, Quinn and Stong to all compact surfaces in the presence of algebraically dual spheres.

Recall from Definition 1.4 that, for $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ as in Convention 1.1, by definition the *Kervaire–Milnor invariant* $\text{km}(F) \in \mathbb{Z}/2$ vanishes if and only if, after finitely many finger moves on the interior of F , taking F to some F' , there is a convenient collection of Whitney discs, with interiors disjoint from F' , pairing all the double points of F' .

The finger moves in this definition are relevant because finger moves can add relations to the fundamental group $\pi_1(M \setminus F)$, making it easier to find (Whitney) discs in the complement of F .

We could have allowed arbitrary regular homotopies, from F to F' , in the definition of km . However, this is not needed, as the following result shows. Note that a nonregular homotopy can change $\text{km}(F)$; see Corollary 6.3.

Proposition 3.1 *Let Σ and M be as in Convention 1.1. If $F_1, F_2: \Sigma \looparrowright M$ are regularly homotopic generic immersions, then $\text{km}(F_1) = \text{km}(F_2) \in \mathbb{Z}/2$.*

Proof To show that $\text{km}(F_1) = \text{km}(F_2)$, by symmetry it suffices to show that $\text{km}(F_1) = 0$ implies $\text{km}(F_2) = 0$. Suppose that $\text{km}(F_1) = 0$, and let F'_1 be obtained from F_1 by finger moves such that the intersections of F'_1 can be paired up by Whitney discs $\{W_i\}$ as in Definition 1.4. Since F'_1 and F_2 are regularly homotopic, there is a generic immersion F_3 such that F_3 can be obtained from both F'_1 and F_2 by finger moves and ambient isotopies. Since F_3 is obtained from F'_1 by finger moves and ambient isotopies and the finger moves can be assumed to be disjoint from $\{W_i\}$, all the double points of F_3 can also be paired up by Whitney discs with interiors disjoint from F_3 , as in Definition 1.4. Since F_3 is obtained from F_2 by finger moves, by taking $F'_2 := F_3$ it follows that $\text{km}(F_2) = 0$. \square

Definition 1.4 is optimised for the proof of Theorem 1.2, as we will see shortly, but is difficult to use in practice. In particular, while one may fortuitously detect specific finger moves and Whitney discs to show $\text{km}(F) = 0$, without a combinatorial description it appears, for a given F , to be hard to prove that the required finger moves from F to some F' , together with Whitney discs for F' , do not exist. We provide precisely such a combinatorial reformulation in Theorem 1.9, generalising Matsumoto's invariant to our formula for $t(F)$ for b -characteristic F . In the proof of Theorem 1.6 we will show that in the presence of dual spheres this agrees with Definition 1.4.

4 The proof of the surface embedding theorem

The surface embedding theorem (Theorem 1.2) can be deduced using the proof of [Freedman and Quinn 1990, Theorem 10.5(1)], combined with an observation in [Powell et al. 2020, Theorem A and Lemma 6.5] for the condition on the homotopy class of \bar{G} , using our definition of the Kervaire–Milnor

invariant (Definition 1.4). Since the surface embedding theorem does not follow directly from the statement of [Freedman and Quinn 1990, Theorem 10.5(1)], as previously discussed, and also since the latter requires a correction by Stong [1994], it can be hard for the uninitiated to piece together a correct proof. Therefore we provide one in this section.

Further, the statement of [Freedman and Quinn 1990, Theorem 10.5(1)] is itself quite complicated, and our version, focussed on the surface embedding problem, may be useful for those looking to apply the technology of Freedman–Quinn without delving into the details.

4.1 Ingredients

The statement of the surface embedding theorem uses the notions of algebraically and geometrically dual spheres. We recall the definitions.

Definition 4.1 Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1, with components $\{f_i\}_{i=1}^m$.

- (1) A collection $G = \{g_i: S^2 \looparrowright M\}_{i=1}^m$ of generic immersions is said to be *algebraically dual* to F if $F \sqcup G$ is a generic immersion and $\lambda(f_i, g_j) = [\delta_{ij}] \in \Gamma_{f_i, g_j}$ for all i and j for some choice of basings for F and G .
- (2) A collection $\bar{G} = \{\bar{g}_i: S^2 \looparrowright M\}_{i=1}^m$ of generic immersions is *geometrically dual* to F if $F \sqcup \bar{G}$ is a generic immersion and the geometric count of intersections satisfies $|f_i \pitchfork \bar{g}_j| = \delta_{ij}$ for all i and j .

We will need the following lemma, the idea behind which is due to Casson [1986]; see also [Freedman 1982, Section 3]. The formulation we give here is from [Powell et al. 2020, Lemma 5.1].

Lemma 4.2 (Geometric Casson lemma) *Let F and G be transversely intersecting generic immersions of compact surfaces in a connected 4–manifold M . Assume that the intersection points $\{p, q\} \subseteq F \pitchfork G$ are paired by a Whitney disc W . Then there is a regular homotopy from $F \cup G$ to $\bar{F} \cup \bar{G}$ such that $\bar{F} \pitchfork \bar{G} = (F \pitchfork G) \setminus \{p, q\}$. That is, the two paired intersections have been removed. The regular homotopy may create many new self-intersections of F and G ; however, these are algebraically cancelling.*

The proof of the surface embedding theorem also relies on Freedman’s disc embedding theorem, whose statement we recall.

Theorem 4.3 (Disc embedding theorem [Freedman 1982; Freedman and Quinn 1990; Powell et al. 2020]; see also [Behrens et al. 2021]) *Let M be a connected topological 4–manifold with good fundamental group, and let*

$$F = \{f_i\}_{i=1}^m: (D^2 \sqcup \dots \sqcup D^2, S^1 \sqcup \dots \sqcup S^1) \looparrowright (M, \partial M)$$

be a generic immersion of finitely many discs. Assume that F has framed algebraically dual spheres $G = \{g_i\}_{i=1}^m \subseteq \pi_2(M)$ such that $\lambda(g_i, g_j) = 0 = \mu(g_i)$ for all $i \neq j$. Then there is a flat embedding

$\bar{F} = \{\bar{f}_i\}_{i=1}^m : (\bigsqcup D^2, \bigsqcup S^1) \hookrightarrow (M, \partial M)$, which is equipped with geometrically dual spheres $\bar{G} = \{\bar{g}_i\}_{i=1}^m$, such that \bar{F} and F have the same framed boundary and $[\bar{g}_i] = [g_i] \in \pi_2(M)$ for all i .

We will also freely use standard constructions such as symmetric contraction, boundary twisting and interior twisting (ie adding local cusps). See [Freedman and Quinn 1990, Chapters 1–2; Powell and Ray 2021a] for further details.

4.2 Proof of the surface embedding theorem

We recall the statement for the convenience of the reader.

Theorem 1.2 (Surface embedding theorem) *Let $F = \{f_i\}_{i=1}^m : (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. Suppose that $\pi_1(M)$ is good and that F has algebraically dual spheres $G = \{g_i\}_{i=1}^m \subseteq \pi_2(M)$. Then the following statements are equivalent:*

- (i) *The self-intersection number $\mu(F)$ and the Kervaire–Milnor invariant $\text{km}(F) \in \mathbb{Z}/2$ vanish.*
- (ii) *There is an embedding $\bar{F} = \{\bar{f}_i\}_{i=1}^m : (\Sigma, \partial\Sigma) \hookrightarrow (M, \partial M)$, regularly homotopic to F relative to $\partial\Sigma$, with geometrically dual spheres $\bar{G} = \{\bar{g}_i : S^2 \looparrowright M\}_{i=1}^m$ such that $[\bar{g}_i] = [g_i] \in \pi_2(M)$ for all i .*

Proof The direction (ii) \implies (i) follows from the fact that the intersection and self-intersection numbers, as well as the Kervaire–Milnor invariant, are invariant under regular homotopy (relative to the boundary) by Propositions 2.18, 2.25 and 3.1, respectively.

The proof of the direction (i) \implies (ii) reduces to the disc embedding theorem (Theorem 4.3) as follows. The argument is similar to the proof of [Freedman and Quinn 1990, Corollary 5.1B] (see also the proof of [Powell et al. 2020, Theorem 8.1]).

Apply the geometric Casson lemma (Lemma 4.2) to upgrade $G = \{g_i\}$ from algebraically to geometrically dual spheres $G' = \{g'_i\}$, changing F to F' by a regular homotopy in the process. The intersection and self-intersection numbers, and the Kervaire–Milnor invariant, vanish for F . So they also vanish for F' , since all three quantities are preserved under regular homotopy relative to the boundary by Propositions 2.18, 2.25 and 3.1.

Then, by the definition of the Kervaire–Milnor invariant (Definition 1.4), after further finger moves changing F' to some F'' , we can find a convenient collection of Whitney discs $\mathcal{W} = \{W_j\}$ for F'' whose interiors are disjoint from F'' . Moreover, F'' and G' are still geometrically dual, since the finger moves may be assumed to miss G' .

We shall apply the disc embedding theorem (Theorem 4.3) to the collection of generically immersed discs \mathcal{W} in the 4-manifold $M \setminus \nu F''$, so we verify that the hypotheses are satisfied. The Whitney discs \mathcal{W} have framed algebraically dual spheres as follows. The Clifford tori at the double points of F'' are geometrically dual to \mathcal{W} . Symmetrically contract half of these tori, one per Whitney disc, using meridional discs for F'' tubed into the geometrically dual spheres $G' = \{g'_i\}$. The resulting spheres

are only algebraically dual to \mathcal{W} since the components of \mathcal{W} and G' may intersect arbitrarily; however, they have vanishing intersection and self-intersection numbers since they were produced by symmetric contraction. They are also framed, as we argue briefly now. If a sphere g_i in G' is not framed, then the symmetric contraction uses incorrectly framed caps. However, in the symmetric contraction process, each cap is used twice, with opposite orientations, and so any framing discrepancies cancel out. Since F'' has geometrically dual spheres, $\pi_1(M \setminus \nu F'') \cong \pi_1(M)$ and is thus good. This verifies the hypotheses of the disc embedding theorem (Theorem 4.3) for \mathcal{W} , as desired.

Apply the disc embedding theorem to the Whitney discs \mathcal{W} in $M \setminus \nu F''$ to obtain disjointly embedded, flat, framed Whitney discs $\{\overline{W}_l\}$ for the double points of F'' , with interiors still disjoint from F'' , along with a collection of geometrically dual spheres for the $\{\overline{W}_l\}$ in $M \setminus \nu F''$.

Tube any intersections of G' with $\{\overline{W}_l\}$ into the geometrically dual spheres for $\{\overline{W}_l\}$, giving a new collection of spheres $\overline{G} = \{\overline{g}_i\}$ disjoint from $\{\overline{W}_l\}$. Now we have that the interiors of $\{\overline{W}_l\}$ lie in the complement of $F'' \cup \overline{G}$, and moreover F'' and \overline{G} are geometrically dual. Perform Whitney moves on F'' along $\{\overline{W}_l\}$ to arrive at an embedding \overline{F} , as claimed. By construction, \overline{F} and \overline{G} are geometrically dual. That $[\overline{g}_i] = [g_i] \in \pi_2(M)$ for each i follows from [Powell et al. 2020, Lemma 6.5]. \square

4.3 The π_1 -negligible surface embedding theorem

Recall that a map $F: X \rightarrow Y$ is called π_1 -negligible if the inclusion $Y \setminus F(X) \subseteq Y$ induces an isomorphism on π_1 for all basepoints. Here is a reformulation of the surface embedding theorem.

Corollary 4.4 (The π_1 -negligible surface embedding theorem) *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. Suppose that F is π_1 -negligible and that $\pi_1(M)$ is good. Then $\mu(F)$ and the Kervaire–Milnor invariant of F both vanish if and only if there exists a π_1 -negligible embedding $\overline{F}: (\Sigma, \partial\Sigma) \hookrightarrow (M, \partial M)$ regularly homotopic to F , relative to the boundary.*

This corollary follows from the surface embedding theorem and the fact that a generic immersion $F: \Sigma \looparrowright M$ is π_1 -negligible if and only if F admits geometrically dual spheres, which can be seen as follows. For the forward direction, the meridional circles are null-homotopic in M , so by π_1 -negligibility they are null-homotopic in $M \setminus F(\Sigma)$. The union of null-homotopies with meridional discs gives geometrically dual spheres. For the reverse direction, first note that, by general position, the homomorphism $\pi_1(M \setminus F(\Sigma)) \rightarrow \pi_1(M)$ is surjective. The kernel is normally generated by a collection consisting of one meridional circle for each connected component of Σ . Since geometrically dual spheres provide null-homotopies for these meridians, the assertion follows. By the geometric Casson lemma (Lemma 4.2), the map F in the statement of the surface embedding theorem (Theorem 1.2) is regularly homotopic to a π_1 -negligible map, due to the existence of the algebraically dual spheres G . Indeed, this is the first step of the proof of the surface embedding theorem. Note that Theorem 1.2 also controls the homotopy class of the dual spheres, and so is slightly stronger than Corollary 4.4.

5 Band characteristic maps and the combinatorial formula

In this section we define b -characteristic surfaces (Definition 5.17) and motivate the combinatorial formula for the Kervaire–Milnor invariant (Definition 1.5). We postpone many of the proofs to Section 7. We hope this will help the reader to assimilate the overall structure more easily. We work towards the definition of b -characteristic surfaces by first defining the related notions of s -characteristic and r -characteristic surfaces, mirroring the historical development. These latter definitions are simpler to state and serve to motivate the more complicated definition of b -characteristic surfaces.

A sphere $g: S^2 \looparrowright M$ in a topological 4-manifold M is said to be *twisted* if the Euler number of the normal bundle is odd. We say g is *framed* if the normal bundle is trivial. To coincide with the usual meaning of framed, one can also implicitly choose a trivialisation, although we will not make use of such a choice. Observe that, if the normal bundle of g has even Euler number, then it is homotopic to a generically immersed sphere with trivial normal bundle, via adding local cusps.

Definition 5.1 Let $F = \{f_1, \dots, f_m\}: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. We define $\Sigma^\circ \subseteq \Sigma$ so that, for $i \in \{1, \dots, m\}$, the component $\Sigma_i \subseteq \Sigma^\circ$ if and only if there is no framed immersed sphere g_i with $\lambda(f_j, g_i) = \delta_{ij}$ for all $j = 1, \dots, m$. Then we use

$$F^\circ = \{f_i^\circ: \Sigma_i^\circ \rightarrow M\}$$

to denote the restriction of F to Σ° . Note that, if an f_i does not admit an algebraically dual sphere at all, then it belongs to F° .

Recall that $x \in H_2(M, \partial M; \mathbb{Z}/2)$ is said to be *characteristic* if $x \cdot a = a \cdot a \in \mathbb{Z}/2$ for every $a \in H_2(M; \mathbb{Z}/2)$, where $-\cdot-$ denotes the intersection pairing $H_2(M; \mathbb{Z}/2) \times H_2(M, \partial M; \mathbb{Z}/2) \rightarrow \mathbb{Z}/2$. The next definition gives a weaker notion.

Definition 5.2 Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. The map F is called *spherically characteristic* (or *s-characteristic* for short) if $F \cdot a = a \cdot a \in \mathbb{Z}/2$ for all $a \in \pi_2(M)$, considered as an element of $H_2(M; \mathbb{Z}/2)$.

We will show in Lemma 5.18 that b -characteristic maps are s -characteristic.

Lemma 5.3 Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1.

- (i) If F is s -characteristic, then $F = F^\circ$.
- (ii) If F has algebraically dual spheres, then F° is s -characteristic or empty.

Proof To prove (i), suppose F is not equal to F° ; then there exists a component Σ_i of Σ with a framed dual sphere g_i , ie with $\lambda(f_j, g_i) = \delta_{ij}$ for all $j \neq i$. This leads to the contradiction

$$1 = f_i \cdot g_i = f_i \cdot g_i + \sum_{j \neq i} f_j \cdot g_i = F \cdot g_i = g_i \cdot g_i = 0 \in \mathbb{Z}/2,$$

where the second-to-last equality follows from F being s -characteristic.

To prove (ii), suppose that F has algebraically dual spheres. Note that the dual spheres for $F^\infty \subseteq F$ are necessarily twisted. Assume that F^∞ neither s -characteristic nor empty. Then there exists $a \in \pi_2(M)$ such that $F^\infty \cdot a \neq a \cdot a \in \mathbb{Z}/2$. By tubing into a dual sphere to a component of F^∞ if necessary, we can assume that a is untwisted, ie $a \cdot a$ is zero and that $F^\infty \cdot a = 1$. Choose some component f_j of F^∞ such that $a \pitchfork f_j$ is nonempty. Except for one of the intersections between f_j and a , tube all the intersections of a and F^∞ into the corresponding dual spheres to F^∞ . Call the resulting sphere a'_j . Since $F^\infty \cdot a = 1$, we tubed into an even number of dual spheres, so $a'_j \cdot a'_j = a \cdot a = 0 \in \mathbb{Z}/2$. Via adding local cusps, we may assume that a'_j is framed. We also have that $\lambda(f_i, a'_j) = \delta_{ij}$ for all i and j . This contradicts the definition of F^∞ . \square

Recall that a *convenient* collection of Whitney discs \mathcal{W} for the intersections within F consists of framed, generically immersed Whitney discs with interiors transverse to F , and with disjointly embedded boundaries. Recall the invariant t from Definition 1.5 appearing in Theorem 1.6, where $t(F, \mathcal{W})$ is the mod 2 count of transverse intersections between F and the interiors of the Whitney discs in \mathcal{W} .

If F is not s -characteristic, then we can change $t(F, \mathcal{W})$ as follows. Given $a \in \pi_2(M)$ with $F \cdot a$ odd but $a \cdot a$ even, one can tube a framed Whitney disc W for F into a framed representative $\tilde{a}: S^2 \looparrowright M$, keeping the new Whitney disc framed but adding an odd number of interior intersections with F . If $F \cdot a$ is even and $a \cdot a$ is odd, one can tube W into a representative \tilde{a} and also add an odd number of boundary twists to keep the new Whitney disc framed but again adding an odd number of interior intersections with F . Using Lemma 5.3, this is one reason for the appearance of F^∞ in the following statements.

The following lemma is also used in the proof of Theorem 1.6, and shows that the vanishing of t for a given collection of Whitney discs implies the vanishing of km .

Lemma 5.4 *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. Suppose that F admits algebraically dual spheres, and that all double points of F are paired by a convenient collection \mathcal{W} of Whitney discs. Let $\mathcal{W}^\infty \subseteq \mathcal{W}$ denote the subcollection of Whitney discs for the intersections within F^∞ , where F^∞ is as in Definition 5.1. If $t(F^\infty, \mathcal{W}^\infty) = 0$, then $\text{km}(F) = 0$.*

We wish to find practically verifiable conditions on F that guarantee that $t(F^\infty, \mathcal{W}^\infty)$ is independent of the collection of Whitney discs \mathcal{W}^∞ . More precisely, the value of $t(F^\infty, \mathcal{W}^\infty)$ should be independent of the pairing of double points, the Whitney arcs joining the paired double points (which includes the choice of sheets at each double point), and finally the Whitney discs. In the case that each Σ_i is simply connected, $t(F^\infty, \mathcal{W}^\infty)$ agrees with [Freedman and Quinn 1990, Definition 10.8A]. However, Freedman and Quinn claim in their Lemma 10.8B that, for simply connected Σ , the quantity $t(F^\infty, \mathcal{W}^\infty)$ only depends on F^∞ , and not on the Whitney discs, as long as F^∞ is s -characteristic. This is not true in general, as pointed out and corrected by Stong [1994]. Further, again with $\pi_1(\Sigma_i) = 1$ for all i , Stong established that the value of t does not depend on the choice of \mathcal{W}^∞ using the notion of r -characteristic discs and spheres. Here is our generalisation of his notion.

Definition 5.5 Let Σ and M be as in Convention 1.1. A map $F: (\Sigma, \partial\Sigma) \rightarrow (M, \partial M)$ is called $\mathbb{R}P^2$ -characteristic (or simply r -characteristic) if $F \cdot R = R \cdot R \in \mathbb{Z}/2$ for every map $R: \mathbb{R}P^2 \rightarrow M$ satisfying $R^*w_1(M) = 0$.

Remark 5.6 A map $c: \mathbb{R}P^2 \rightarrow S^2$ of odd degree (eg a collapse map) composed with elements of $\pi_2(M)$ can be used to show that r -characteristic maps are s -characteristic. Indeed, given F and $a \in \pi_2(M)$, we obtain $a \circ c: \mathbb{R}P^2 \rightarrow M$, and we have $0 = F \cdot (a \circ c) + (a \circ c) \cdot (a \circ c) = F \cdot a + a \cdot a \in \mathbb{Z}/2$, where the second equality uses that c has odd degree.

Stong’s key observation [1994] was that, in some instances involving elements of order two in $\pi_1(M)$, one can change the choice of sheets at two double points of F° , and hence the Whitney arcs and corresponding Whitney disc, with a resulting change in the value of $t(F^\circ, \mathcal{W}^\circ)$ by one. The restriction to r -characteristic maps removes this source of indeterminacy. To summarise, Stong showed the following theorem.

Theorem 5.7 [Stong 1994] *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1, with Σ a union of discs or spheres. Suppose $\mu(F) = 0$ and that F admits algebraically dual spheres. If F° is not r -characteristic, then $\text{km}(F) = 0$, and if F° is r -characteristic, then $\text{km}(F) = t(F^\circ, \mathcal{W}^\circ)$ for any choice of Whitney discs \mathcal{W}° for the intersections within F° .*

Remark 5.8 Combining Theorems 2.32 and 5.7 with Theorem 1.2 gives the complete answer to the embedding problem for spheres and discs with algebraically dual spheres for good fundamental groups, due to Freedman, Quinn and Stong. Theorem 2.32 allows one to fix the regular homotopy class of generic immersions within the homotopy class to be that with $\mu(-)_1 = 0$, Theorem 5.7 computes the Kervaire–Milnor invariant, and then one applies Theorem 1.2 to conclude whether or not there is a regular homotopy to an embedding.

Our contribution in the present paper extends this solution to the case that the components of Σ are not all simply connected. In this case there is a further source of indeterminacy coming from the choice of Whitney arcs on Σ . For this reason we need a stronger restriction on F° .

Definition 5.9 A *band* refers to either of the two D^1 -bundles over S^1 , ie a band is either an annulus or a Möbius band. Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. Let \mathcal{M}_F be the mapping cylinder of F . Write $\mathcal{B}(F) \subseteq H_2(M, \Sigma; \mathbb{Z}/2) := H_2(\mathcal{M}_F, \Sigma; \mathbb{Z}/2)$ for the subset of elements of the relative homology group that can be represented by a square

$$\begin{array}{ccc} \partial B & \xrightarrow{h} & \Sigma \\ \downarrow \iota & & \downarrow F \\ B & \xrightarrow{g} & M \end{array}$$

where B is a band and $\iota: \partial B \hookrightarrow B$ is the inclusion such that

$$(5-1) \quad \langle w_1(M), g(C) \rangle + \langle w_1(\Sigma), h(\partial B) \rangle = 0 \in \mathbb{Z}/2,$$

where C is the core curve of B .




Figure 5: An annular band B (blue) is shown with boundary on F (black). One of the boundary components of B is nonorientable on F and one is orientable, so $\langle w_1(\Sigma), h(\partial B) \rangle = 1$. Therefore, in order for this to be an element of $\mathcal{B}(F)$, we must have $\langle w_1(M), g(C) \rangle = 1$, where M is the ambient 4-manifold and C is the core curve of the annulus, shown in blue.

See Figure 5 for an example of a band. Note that every element of $\mathcal{B}(F)$ can be represented by a generic immersion of pairs $(B, \partial B) \looparrowright (M, \Sigma)$ (Definition 2.6). Writing $H_2(M, \Sigma; \mathbb{Z}/2)$ in place of $H_2(\mathcal{M}_F, \Sigma; \mathbb{Z}/2)$ is a slight but standard abuse of notation. The pair (g, h) induces a relative homology class since the map

$$g \sqcup (h \times \text{Id}_{[0,1]}): B \sqcup (\partial B \times [0, 1]) \rightarrow M \sqcup (\Sigma \times [0, 1])$$

descends to a map $\mathcal{M}_i \rightarrow \mathcal{M}_F$. The mapping cylinder \mathcal{M}_i is homeomorphic to B , and so we obtain a map $(B, \partial B) \rightarrow (\mathcal{M}_F, \Sigma)$. The image of the relative fundamental class $[B, \partial B]$ in $H_2(\mathcal{M}_F, \Sigma; \mathbb{Z}/2)$ is an element of $\mathcal{B}(F)$. From now on, since $\mathcal{M}_F \simeq M$, to simplify the notation we will not mention the mapping cylinder and refer to $\mathcal{B}(F) \subseteq H_2(M, \Sigma; \mathbb{Z}/2)$.




Figure 6: (a) An immersed surface S (blue) in the ambient manifold M , and the image (black) of a generic immersion $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$. (b) A thin tube is added, with one boundary component on Σ and one on S . The surface B (blue) is obtained by cutting out a disc on S and gluing in the tube. Note that, compared to S , the surface B has a new boundary component lying on Σ .

We will see in Lemma 7.7 that, given a generic immersion $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ as in Convention 1.1, every element of $H_2(M, \Sigma; \mathbb{Z}/2)$ can be represented by an immersion of some compact surface S into M , with interior transverse to F , and with boundary generically immersed in $F(\Sigma)$ away from the double points. The subset $\mathcal{B}(F)$ consists of those homology classes for which S can be chosen to be a band satisfying condition (5-1).

We use the notation

$$\partial: H_2(M, \Sigma; \mathbb{Z}/2) \rightarrow H_1(\Sigma; \mathbb{Z}/2)$$

for the connecting homomorphism from the long exact sequence of the pair. A class represented by a compact surface S is mapped to its boundary ∂S under the map ∂ . Then $\partial\mathcal{B}(F) \subseteq H_1(\Sigma; \mathbb{Z}/2)$ consists of (the homology classes of) those closed 1-manifolds immersed in Σ whose images under F bound bands in M satisfying (5-1).

Construction 5.10 Given a generic immersion $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ as in Convention 1.1, suppose we have a generically immersed surface S in M with boundary on Σ , ie admitting maps satisfying

$$\begin{array}{ccc} \partial S & \xrightarrow{h} & \Sigma \\ \downarrow \iota & & \downarrow F \\ S & \xrightarrow{g} & M \end{array}$$

where possibly ∂S is empty. Then the tubing procedure shown in Figure 6 can be used to create a band, as follows. If S is a disc, the procedure gives an annulus B with boundary lying on Σ . This annulus satisfies (5-1), and therefore lies in $\mathcal{B}(F)$, if and only if $\langle w_1(\Sigma), h(\partial S) \rangle = 0 \in \mathbb{Z}/2$, since the core of B is null-homotopic in M and the newly created boundary component of B is null-homotopic in Σ .

In the case that S is a sphere, we can perform the tubing procedure of Figure 6 to S twice. In this case, both boundary components of the annulus created are null-homotopic on Σ , so we always produce an element of $\mathcal{B}(F)$.

Finally, if S is an $\mathbb{R}P^2$, the tubing procedure creates a Möbius band with boundary on Σ , which lies in $\mathcal{B}(F)$ if and only if $\langle w_1(M), g(\mathbb{R}P^1) \rangle = 0 \in \mathbb{Z}/2$, where $\mathbb{R}P^1 \subseteq \mathbb{R}P^2$.

When defining b -characteristic surfaces, we will restrict to the case that the $\mathbb{Z}/2$ -valued intersection form λ_Σ is trivial on $\partial\mathcal{B}(F)$. We can restrict in this way because of the following lemma.

Lemma 5.11 Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1, with $\mu(F) = 0$. If the $\mathbb{Z}/2$ -valued intersection form λ_Σ on $H_1(\Sigma; \mathbb{Z}/2)$ is nontrivial on $\partial\mathcal{B}(F)$, then we can change F by a regular homotopy to F' such that there are convenient collections of Whitney discs \mathcal{W} and \mathcal{W}' for the double points of F and F' , respectively, such that $t(F, \mathcal{W}) \neq t(F', \mathcal{W}')$.

Moreover, if F has dual spheres and the $\mathbb{Z}/2$ -valued intersection form λ_{Σ^∞} on $H_1(\Sigma^\infty; \mathbb{Z}/2)$ is nontrivial on $\partial\mathcal{B}(F^\infty)$ then $\text{km}(F) = 0$.

In the case that $\lambda_\Sigma|_{\partial\mathcal{B}(F)}$ is trivial, we define an invariant Θ on the set $\mathcal{B}(F)$. We will need the following notions:

- (1) For Z a closed 1–manifold generically immersed in Σ , the self-intersection number $\mu_\Sigma(Z) \in \mathbb{Z}/2$ of Z counts the number of double points, which we assume without loss of generality to be disjoint from the double points of F . As usual, this is not invariant under homotopies of Z in Σ , only under regular homotopies.
- (2) Let S be a compact surface, with a generic immersion of pairs $(S, \partial S) \looparrowright (M, \Sigma)$. Suppose that $w_1(\Sigma)$ is trivial on each component of ∂S , eg if Σ is orientable. Then the normal bundle of ∂S in Σ is trivial and we can pick a nowhere-vanishing section (if S is closed, this is an empty choice). Extend this section to the normal bundle of S in M (such a normal bundle exists by Corollary 2.7) such that the extension is transverse to the zero section. Then we define the *Euler number* $e(S)$ to be the number of zeros of this section modulo 2. Observe that this is analogous to how one measures the twisting of a Whitney disc with respect to the Whitney framing. For S a closed surface, this coincides with the usual definition of the Euler number.

Here is the definition of $\Theta(S)$ in the case that $w_1(\Sigma)$ is trivial on every component of ∂S .

Definition 5.12 Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1, with $\mu(F) = 0$. Let A be a choice of Whitney arcs pairing the double points of F . Let S be a compact surface in M with a generic immersion of pairs $(S, \partial S) \looparrowright (M, \Sigma)$ such that ∂S is transverse to A and $w_1(\Sigma)$ is trivial on every component of ∂S . Define

$$(5-2) \quad \Theta_A(S) := \mu_\Sigma(\partial S) + |\partial S \pitchfork A| + |\text{Int } S \pitchfork F| + e(S) \pmod{2}.$$

For closed S we have $\Theta_A(S) \equiv |\text{Int } S \pitchfork F| + e(S) \pmod{2}$, and thus $\Theta_A(S)$ vanishes for all closed S if and only if F is characteristic in the traditional sense. In the proof of Theorem 1.6, we will only use the definition of Θ_A for bands. But the case of general surfaces will be useful for our proof that, in the cases relevant to us, Θ_A does not depend on the choice of A (see Lemma 5.16 below).

Remark 5.13 Definition 5.12 suffices in the case of orientable Σ . The reader only interested in this case may safely skip ahead to Lemma 5.16.

If a component of ∂S is orientation-reversing in Σ , then its normal bundle in Σ is nontrivial and hence we may not use it to choose a nowhere-vanishing section of the normal bundle of S on its boundary as before to define the Euler number. However, bands with such boundaries may exist in the ambient 4–manifold and must be considered. When $w_1(\Sigma)$ is nontrivial on precisely one component of ∂S , eg when ∂S is connected, we know $\lambda_\Sigma(\partial S, \partial S) = 1$. Then, by Lemma 5.11, if a band with such boundary exists, then $\text{km}(F) = 0$, and there is no need to define Θ . In particular, note that this means that we need not consider the case of a Möbius band whose boundary consists of a curve on which $w_1(\Sigma)$ is nontrivial. There is one final case of relevant bands left to consider, which we do next.




Figure 7: (a) A band B (blue) is shown for F (black). Both boundary components of B are nonorientable on Σ . A vertical arc D is shown in orange. (b) By cutting along the arc D we obtain a disc \hat{B} . We show how to choose a nowhere-vanishing section of the normal bundle of $\partial\hat{B}$ in M . (c) Adding a tube to F guided by D splits the band into a disc Δ , and changes Σ to a surface Σ' . We show how to choose a section of the normal bundle of $\partial\Delta$ in F' .

Definition 5.14 Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1, with $\mu(F) = 0$. Let A be a choice of Whitney arcs pairing the double points of F . Let B be an annulus with a generic immersion of pairs $(B, \partial B) \looparrowright (M, \Sigma)$, such that ∂B transverse to A and $w_1(\Sigma)$ is nontrivial on both components of ∂B . Pick an embedded arc D in B connecting the two boundary components, disjoint from the self-intersections of B and the intersections of B with F . Let \hat{B} be the result of cutting B open along D . There is a canonical quotient map $\gamma: \hat{B} \rightarrow B$.

Let ν_B^M denote the normal bundle of B in M . Pick a nowhere-vanishing section of $\gamma^*\nu_B^M$ on $\partial\hat{B}$ as follows. On each part of $\partial\hat{B}$ that maps to ∂B , use the normal bundle of ∂B in F to define the section locally. For this we require that, on ∂D , the two vectors for the two components agree up to multiplication by ± 1 . On the part of $\partial\hat{B}$ that maps to D , pick a section so that, on every pair of points that map to the same point in D , the vectors agree up to multiplication by ± 1 . See Figure 7(b). We define

$$(5-3) \quad \Theta_A(B, D) := \mu_\Sigma(\partial B) + |\partial B \pitchfork A| + |\text{Int } B \pitchfork F| + e(\hat{B}) \pmod 2.$$

Remark 5.15 An alternative definition of $\Theta_A(B, D)$ would use the arc D to add a tube to $F(\Sigma)$ in such a way that the tube intersects the band B in two parallel copies of D . More precisely, we perform an ambient surgery on $F(\Sigma)$. This requires choosing a 2-dimensional subbundle of the normal bundle of D in M — the tube itself consists of the circle bundle for this subbundle, considered within a tubular neighbourhood of D . We build the required subbundle by first choosing a section lying in the normal bundle of D in B , denoted by ν_D^B . The second section can be chosen freely. Let F' denote the immersion constructed by the tubing procedure. Observe that the domain of F' , denoted by Σ' , is obtained from the abstract surface Σ by adding a 1-handle. Depending on the choice of the second section above, this may be an orientation-reversing or orientation-preserving 1-handle, but this will not matter for us.

Adding the tube changes the band B to a disc Δ , by removing a thin strip neighbourhood of D . The disc Δ has boundary lying on Σ' . Observe that $\partial\Delta$ is orientation-preserving on Σ' , since it is the result of banding together two orientation-reversing curves. Since no new intersection points were added, the collection A is a collection of Whitney arcs for the intersection points of F' . As a result, we may define

$$\Theta_A(B, D) := \Theta_A(\Delta),$$

where the latter is computed as in Definition 5.12. In order to see that this agrees with Definition 5.14, we need only check that the definition of the Euler number terms agree, assuming we choose the tube to be thin enough to miss any double points. For this compare Figure 7(b)–(c) to see that the sections at the boundary in both cases are the same and hence also the Euler numbers coincide. When comparing the pictures, the choice of the second section of the 2–dimensional subbundle which determines the tube corresponds to the choice of the section of $\gamma^* \nu_B^M$ on the part of $\partial\hat{B}$ that maps to D .

Note that we did not prove that $e(\hat{B})$ is independent of the choice of D . This will follow from the upcoming proof of Lemma 5.16(ii) below, which states that $\Theta_A(B, D)$ depends only on the homology class of B . The following lemma, whose proof is again deferred to Section 7, shows that Θ is well defined in all required cases. As a reminder, the case of orientable Σ does not require the notion of $\Theta_A(B, D)$ from Definition 5.14, so parts (ii) and (iii) of the following lemma may be skipped by anyone only interested in that case.

Lemma 5.16 *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1, with $\mu(F) = 0$. Let A be a choice of Whitney arcs pairing the double points of F .*

- (i) *Let S be a compact surface with a generic immersion of pairs $(S, \partial S) \looparrowright (M, \Sigma)$ such that ∂S is transverse to A and $w_1(\Sigma)$ is trivial on every component of ∂S . Then $\Theta_A(S) \in \mathbb{Z}/2$ depends only on the homology class of S in $H_2(M, \Sigma; \mathbb{Z}/2)$.*
- (ii) *Let B be an annulus with a generic immersion of pairs $(B, \partial B) \looparrowright (M, \Sigma)$ such that ∂B is transverse to A and $w_1(\Sigma)$ is nontrivial on both components of ∂B . Pick an embedded arc D in B connecting the components of ∂B and disjoint from all double points. Then $\Theta_A(B, D) \in \mathbb{Z}/2$ depends only on the homology class of B in $H_2(M, \Sigma; \mathbb{Z}/2)$. In particular, $\Theta_A(B, D)$ does not depend on D , so we write $\Theta_A(B)$.*
- (iii) *Let S be a surface as in (i) and let B be an annulus as in (ii) such that $[S] = [B] \in H_2(M, \Sigma; \mathbb{Z}/2)$. Then $\Theta_A(S) = \Theta_A(B) \in \mathbb{Z}/2$.*
- (iv) *If $\lambda_\Sigma|_{\partial\mathfrak{B}(F)} = 0$, the restriction of Θ_A to $\mathfrak{B}(F)$ is independent of the choice of A , giving a well-defined map $\Theta: \mathfrak{B}(F) \rightarrow \mathbb{Z}/2$.*

Finally, we are ready to define the required generalisation of r –characteristic maps, called b –characteristic maps.

Definition 5.17 Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1, with $\mu(F) = 0$. We say F is *band characteristic* (or *b-characteristic* for short) if $\lambda_\Sigma|_{\partial\mathcal{B}(F)} = 0$ and $\Theta: \mathcal{B}(F) \rightarrow \mathbb{Z}/2$ is trivial.

Lemma 5.18 Every *b-characteristic* map is *r-characteristic*. Moreover, the two notions agree for unions of discs or spheres.

Proof Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1, with $\mu(F) = 0$. Let $R: \mathbb{R}P^2 \rightarrow M$ be a generic immersion which is transverse to F and such that $R^*w_1(M) = 0$. We apply Construction 5.10. In other words, take a small disc on $F(\Sigma)$ away from $R \pitchfork F$, and tube into the image of R . This creates a Möbius band B with boundary on Σ . Here ∂B is homotopically trivial in Σ , so $\langle w_1(\Sigma), \partial B \rangle = 0$. The core C of B corresponds to $\mathbb{R}P^1$ within the original immersed $\mathbb{R}P^2$. Therefore, since $R^*w_1(M) = 0$, we have that $\langle w_1(M), C \rangle = 0$. So $B \in \mathcal{B}(F)$. Note that $\Theta: \mathcal{B}(F) \rightarrow \mathbb{Z}/2$ is well defined by Lemma 5.16(iv) since F is *b-characteristic*. Further, $\Theta(B) = \Theta(R) = F \cdot R + R \cdot R \in \mathbb{Z}/2$ by Lemma 5.16(i) since $[B] = [R] \in H_2(M, \Sigma; \mathbb{Z}/2)$. But this vanishes since F is *b-characteristic*. Hence, F is *r-characteristic*.

For the second sentence, suppose that Σ is a union of discs or spheres and is *r-characteristic*. Let $B \in \mathcal{B}(F)$ be a band. Since Σ is simply connected, the boundary of B is null-homotopic in $F(\Sigma)$. Therefore, B can be closed up using a codimension zero submanifold of Σ to either a sphere or an $\mathbb{R}P^2$ immersed in M . The resulting closed surface R again satisfies $\Theta(B) = F \cdot R + R \cdot R \in \mathbb{Z}/2$ by Lemma 5.16(i). Here $\lambda_\Sigma|_{\partial\mathcal{B}(F)} = 0$ since Σ is simply connected and so $\Theta: \mathcal{B}(F) \rightarrow \mathbb{Z}/2$ is again well defined by Lemma 5.16(iv). Once again, since null-homotopic circles on Σ must be orientation-preserving, $\langle w_1(\Sigma), \partial B \rangle = 0$ and so (5-1) implies that $R^*w_1(M) = 0$. So $\Theta(B) = 0$ since F is *r-characteristic*. Thus, F is *b-characteristic*. It follows that the notions of *b-characteristic* and *r-characteristic* coincide, as claimed. □

Recall that, if F^∞ is *b-characteristic*, then Theorem 1.6 states that $\text{km}(F) = t(F^\infty, \mathcal{W}^\infty)$, so we have a combinatorial description of $\text{km}(F)$. Moreover, since Θ only depends on the homology class of a band in $H_2(M, \Sigma; \mathbb{Z}/2)$, we can in principle determine whether or not F is *b-characteristic* by computing Θ on finitely many homology classes. Having said that, as mentioned in the introduction, in practice deciding precisely which homology classes can be represented by maps of bands may be tricky.

Lemma 5.19 Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1, with $\mu(F) = 0$. If G is regularly homotopic to F and F is *b-characteristic*, then G is *b-characteristic*.

Proof By definition, a regular homotopy can be decomposed into a sequence of ambient isotopies, finger moves and Whitney moves. None of these affect which classes of $H_1(\Sigma; \mathbb{Z}/2)$ bound a band. In particular, $\lambda_\Sigma|_{\partial\mathcal{B}(F)} = \lambda_\Sigma|_{\partial\mathcal{B}(G)}$. Assume that G is not *b-characteristic*. Then either $\lambda_\Sigma|_{\partial\mathcal{B}(G)}$ is nontrivial or there is a band in $\mathcal{B}(G)$ on which Θ is nonvanishing. In the former case, $\lambda_\Sigma|_{\partial\mathcal{B}(F)} = \lambda_\Sigma|_{\partial\mathcal{B}(G)}$ implies that F is not *b-characteristic*.

For the latter case, let B in $\mathcal{B}(G)$ be such that $\Theta(B) = 1$. It suffices to show that such a band still exists after an ambient isotopy, a finger move or a Whitney move. This is obvious for ambient isotopy. Recall that Θ only depends on the homology class of the band by Lemma 5.16. Hence, we can assume that the boundary of B is away from the singularity of the finger move. Then we can still consider the band B as a band for the surface after the finger move and Θ is unchanged. The argument in the case of a Whitney move is similar. We first let B undergo a homotopy to arrange that it is disjoint from the boundary of the Whitney disc W along which the Whitney move is performed. Then we can again consider the same band B for the new surface. The Whitney move leaves all terms in the definition of Θ except $|\text{Int } B \cap F|$ unchanged. Since the Whitney move uses two copies of the Whitney disc, the change in $|\text{Int } B \cap F|$ is twice $|\text{Int } B \cap W|$. As Θ takes values in $\mathbb{Z}/2$, $\Theta(B)$ is unchanged, as claimed. Thus, we have a band B in $\mathcal{B}(F)$ with $\Theta(B) = 1$, and so again F is not b -characteristic. We have shown the contrapositive of the desired statement. \square

Remark 5.20 As a counterpoint to Lemma 5.19, there exist maps that are homotopic to each other but where one is b -characteristic and the other is not. For example, let Σ be the Klein bottle. Then an embedding $f: \Sigma \hookrightarrow \mathbb{R}^4$ must have normal Euler number $e(vf) \in \{-4, 0, 4\}$ by [Massey 1969]. It can be verified, as we do presently, that the embeddings with $e(vf) = 0$ are precisely those which are b -characteristic. Hence, the b -characteristic notion is not invariant under homotopy.

To see that f is b characteristic if and only if $e(vf) = 0$, think of $\Sigma \cong \mathbb{R}\mathbb{P}^2 \# \mathbb{R}\mathbb{P}^2$, with a corresponding isomorphism $H_1(\Sigma; \mathbb{Z}/2) \cong \mathbb{Z}/2 \oplus \mathbb{Z}/2$. There is a standard embedding of $\mathbb{R}\mathbb{P}^2$ in \mathbb{R}^4 with normal Euler number ± 2 , and there are essentially three ways to take connected sums of these embeddings, realising the three options $e(vf) \in \{-4, 0, 4\}$. With $e(vf)$ fixed, these embeddings are unique up to regular homotopy by Theorem 2.32.

We explicitly construct the standard embeddings, as follows. Take two disjoint, unlinked, unknotted Möbius bands M_1 and M_2 in \mathbb{R}^3 , with an $\varepsilon_i \in \{\pm 1\}$ signed half-twist for $i = 1, 2$. Take the boundary connected sum $M_1 \natural M_2$ ambiently to obtain a punctured Klein bottle in \mathbb{R}^3 with boundary an unknot. Cap this unknot off with a standard slice disc in $\mathbb{R}_{\geq 0}^4$ to obtain a standard embedding f with normal Euler number $-2(\varepsilon_1 + \varepsilon_2)$. We do not justify the sign, which depends on conventions that are not important for us.

By Lemma 5.19, it suffices to check whether the three standard embeddings above are b -characteristic, which we do next. First one computes that $\lambda_{\Sigma|_{\partial\mathcal{B}(f)}} = 0$ in all three cases, as follows. We have

$$\mathcal{B}(f) \subseteq H_2(\mathbb{R}^4, \Sigma; \mathbb{Z}/2) \xrightarrow{\cong} H_1(\Sigma; \mathbb{Z}/2) \cong \mathbb{Z}/2 \oplus \mathbb{Z}/2.$$

By (5-1), in order to have $B \in \mathcal{B}(f)$, we need $\langle w_1(\Sigma), h(\partial B) \rangle = 0$. Hence, $\partial B \in H_1(\Sigma; \mathbb{Z}/2) \cong \mathbb{Z}/2 \oplus \mathbb{Z}/2$ is either $(0, 0)$ or $(1, 1)$. To see that (x, x) for $x \in \{0, 1\}$ can be realised as the boundary of a band, pick a simple closed curve Z on Σ representing the homology class (x, x) and a generically immersed disc D bounded by Z in \mathbb{R}^4 . Then add a tube from D to Σ to turn D into an annulus B , using Construction 5.10 (see Figure 6). Thus, $\partial\mathcal{B}(f) = \{(0, 0), (1, 1)\}$, on which λ_{Σ} vanishes.

Hence, whether or not the given standard embedding of Σ is b -characteristic is decided by $\Theta(B)$, where B is a band with boundary $(1, 1) \in H_1(\Sigma; \mathbb{Z}/2)$. For the standard embeddings constructed above, such a band can be constructed explicitly, as follows. Take the core curves of M_1 and M_2 , and connect sum them inside $M_1 \natural M_2$. This gives an unknot representing $(1, 1)$, which bounds a standard slice disc D in $\mathbb{R}_{\leq 0}^4$. Construction 5.10 converts D to a band B .

Since Θ only depends on the class of a band in $H_2(M, \Sigma; \mathbb{Z}/2)$, we can use the band B from the previous paragraph. We shall compute that $\Theta(B) = 0$ for this band if and only if $e(vf) = 0$, that is, if the embedding of Σ arises from the connected sum of the standard embeddings of $\mathbb{R}P^2$ with opposite normal Euler numbers. The curve ∂B is orientation-preserving on Σ , so we use (5-2) to compute $\Theta(B)$. Most of the terms in this definition are trivial in this case, since we are working with an embedding of Σ and ∂B is itself embedded. Also D has interior disjoint from the image of f , and therefore so does B . Only the relative Euler number term remains, which can be computed from the twists in M_1 and M_2 . It follows that $\Theta(B) \equiv \frac{1}{2}(\varepsilon_1 + \varepsilon_2) \in \mathbb{Z}/2$, which vanishes if and only if $\varepsilon_1 = -\varepsilon_2$, which in turn holds if and only if $e(vf) = 0$.

6 Homotopy versus regular homotopy

In this short section, we describe Construction 6.1 and apply it to prove Theorem 1.12, which we used in Section 1.4 to compare homotopy and regular homotopy of maps. Note that the results in this section require that the surface Σ from Convention 1.1 is nonorientable.

If we are interested in finding an embedding in a given homotopy class, rather than a regular homotopy class, we may use the construction below to replace a given map by a homotopic map for which the invariant t is trivial. In particular, the construction is applicable in the cases from Theorem 2.32 in which there are infinitely many regular homotopy classes with $\mu(-)_1 = 0$ in a given homotopy class.

Construction 6.1 Let $F = \{f_i\}_{i=1}^m : (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1, with $\mu(F) = 0$. Let \mathcal{W} denote a convenient collection of Whitney discs for the double points of F . Suppose that there exists f_i with $w_1(\Sigma)|_{\ker(f_i)_\bullet}$ nontrivial.

Let $N \subseteq \Sigma_i$ be a Möbius band with $f_i|_N$ π_1 -trivial. In a small disc in N introduce four double points with the same sign by cusp homotopies and call the resulting immersion f'_i . Let F' denote the map given by $\{f_j\}_{j \neq i} \cup \{f'_i\}$. Then there is a convenient collection of Whitney discs \mathcal{W}' for all the double points of F' such that

$$t(F', \mathcal{W}') \equiv t(F, \mathcal{W}) + 1 \pmod{2}.$$

While we have created four double points with the same sign, we will use in the proof that the Möbius band is nonorientable to change the sign of two of the double points, in order to then be able to find new Whitney discs.

Proof We will pair up the two new pairs of double points with new Whitney discs. Pick any pair of the four new double points and pair them by arcs in the small disc, as in Definition 2.26, such that the resulting circle is null-homotopic in M . For one of the arcs perform a connected sum in the interior with the core α of N . With this new pair of arcs, the double points have opposite sign, and by our choice of N the resulting Whitney circle bounds a Whitney disc W_1 in M with embedded boundary. By boundary twisting, arrange that W_1 is framed, and, by pushing off, ensure there is no intersection between the boundary of W_1 and the boundaries of the components of ${}^{\circ}\mathcal{W}$. For this we push the boundary arc of W_1 over the end of the boundary arc for the Whitney disc in ${}^{\circ}\mathcal{W}$. This way $t(F, {}^{\circ}\mathcal{W})$ remains unchanged. Do the same for the remaining two new double points in f'_i , namely pair them by a Whitney disc W_2 , which by definition is a parallel copy of W_1 .

Since $\lambda_N(\alpha, \alpha) = 1$, the boundaries of W_1 and W_2 intersect an odd number of times. To turn ${}^{\circ}\mathcal{W} \cup \{W_1, W_2\}$ into a convenient collection of Whitney discs we have to remove any intersections between their boundaries. For such an intersection, push the Whitney arc of W_1 over the end of the Whitney arc of W_2 . This will in turn change the number of intersections between the interior of W_1 and f'_i by one mod 2; that is, $|\text{Int } W_1 \pitchfork F| \equiv |\text{Int } W_2 \pitchfork F| + 1 \pmod 2$. Let ${}^{\circ}\mathcal{W}'$ be the resulting collection of Whitney discs. We have

$$t(F', {}^{\circ}\mathcal{W}') \equiv t(F, {}^{\circ}\mathcal{W}) + |\text{Int } W_1 \pitchfork F| + |\text{Int } W_2 \pitchfork F| \equiv t(F, {}^{\circ}\mathcal{W}) + 1 \pmod 2. \quad \square$$

With the above construction in hand, we can now prove Theorem 1.12 from the introduction.

Theorem 1.12 *Let $F = \{f_i\}_{i=1}^m : (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1 with $\mu(F) = 0$. Suppose that there is at least one $f_i \in F^{\circ}$ with $w_1(\Sigma)|_{\ker(f_i)}$ nontrivial. Then there exists a generic immersion F' homotopic to F with $\mu(F') = 0$, and a convenient collection of Whitney discs ${}^{\circ}\mathcal{W}'$ such that $t((F')^{\circ}, ({}^{\circ}\mathcal{W}')^{\circ}) = 0$. Thus, if F' has algebraically dual spheres, then $\text{km}(F') = 0$, and if moreover $\pi_1(M)$ is good, then F' is regularly homotopic, relative to $\partial\Sigma$, to an embedding.*

Proof By the vanishing of the intersection and self-intersection numbers, there is a convenient collection of Whitney discs W for F and therefore for F° . If $t(F^{\circ}, W^{\circ}) = 0$, set $F' = F$. If $t(F^{\circ}, W^{\circ}) = 1$, use Construction 6.1 to find a generic immersion F' homotopic to F , with $t((F')^{\circ}, (W')^{\circ}) = 0$. If F' has algebraically dual spheres, then $\text{km}(F') = 0$ by Theorem 1.6 since either F' is not b -characteristic or $\text{km}(F') = t((F')^{\circ}, (W')^{\circ}) = 0$. If in addition $\pi_1(M)$ is good, apply Theorem 1.2 to see that F' is regularly homotopic, relative to $\partial\Sigma$, to an embedding. \square

We end this section by giving another pair of applications of Construction 6.1.

Proposition 6.2 *Let $f : (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be a generic immersion as in Convention 1.1. Assume that Σ is connected, $\mu(f) = 0$ and $w_1(\Sigma)|_{\ker f}$ is nontrivial while $f^*(w_1(M))$ is trivial. If f' is a generic immersion homotopic to f , both f and f' are b -characteristic, and $e(vf) - e(vf') = \pm 8$, then*

$$t(f') \equiv t(f) + 1 \pmod 2.$$

Proof Since $f^*(w_1(M))$ is trivial, regular homotopy classes of generic immersions homotopic to f are detected by the Euler number of the normal bundle by Theorem 2.32. Further, since $e(vf) - e(vf') = \pm 8$, we may add four cusps (of the same sign) to f to obtain a map f'' which is regularly homotopic to f' . The map f is b -characteristic by assumption, while the map f'' is b -characteristic since f' is, by Lemma 5.19. By Lemma 2.24, $\mu(f)_1 \in \mathbb{Z}/2$, so, by construction, $\mu(f) = \mu(f'') = \mu(f') = 0$. So the quantities $t(f)$ and $t(f'')$ are defined and, further, $t(f'') = t(f')$. Apply Construction 6.1 to see that $t(f'') \equiv t(f) + 1 \pmod{2}$. \square

Applying Proposition 6.2 to immersions of $\mathbb{R}P^2$ into \mathbb{R}^4 , we obtain the following corollary, obstructing generic immersions of $\mathbb{R}P^2$ in \mathbb{R}^4 with $e(f) \not\equiv \pm 2 \pmod{16}$ from being regularly homotopic to an embedding and thus partially recovering the result, due to Massey [1969], that every embedding of $\mathbb{R}P^2$ in \mathbb{R}^4 must have Euler number ± 2 . Massey stated the result for smooth embeddings, since he used the G -signature theorem. But the G -signature theorem was later extended to the topological category by [Wall 1970, Chapter 14B], so Massey's result also holds for locally flat embeddings of $\mathbb{R}P^2$ in \mathbb{R}^4 .

Corollary 6.3 *Let $f: \mathbb{R}P^2 \looparrowright \mathbb{R}^4$ be a generic immersion with $\mu(f) = 0$. Then $t(f) = 0$ if and only if $e(vf) = \pm 2 \pmod{16}$.*

Proof Recall that there exist embeddings $g_{\pm}: \mathbb{R}P^2 \hookrightarrow \mathbb{R}^4$ with Euler number ± 2 . First we prove that $e(vf) \equiv 2 \pmod{4}$. By Lemma 2.24, we know that $\mu(f)_1 \in \mathbb{Z}/2$. Then $\mu(g_+) = 0$, and so $\mu(g_+)_1 = 0$. Since f is homotopic to g_+ and $\mu(f) = 0$, it follows from Theorem 2.32 that $e(vf) \equiv e(vg_+) \equiv 2 \pmod{4}$. (The same argument would have applied with g_- .)

Note that any generic immersion of $\mathbb{R}P^2$ into \mathbb{R}^4 , and in particular the map f , is b -characteristic since $H_2(\mathbb{R}^4, \mathbb{R}P^2; \mathbb{Z}/2) \cong \mathbb{Z}/2$ and the nontrivial element does not satisfy condition (5-1) in Definition 5.9.

We have $t(g_{\pm}) = 0$ since t vanishes for embeddings. Since $e(vf) \equiv 2 \pmod{4}$, it differs from one of ± 2 by a multiple of 8. Let $k \in \mathbb{Z}$ be such that $e(vf) = \pm 2 + 8k = e(vg_{\pm}) + 8k$. By Proposition 6.2,

$$t(f) \equiv t(g_{\pm}) + k \equiv k \pmod{2}.$$

Thus, $t(f) = 0$ if and only if k is even, which is the case precisely when $e(vf)$ differs from $e(vg_{\pm}) = \pm 2$ by a multiple of 16. \square

7 Proofs of statements from Section 5

In this section, we provide the proofs we skipped in Section 5. The following transfer move will be useful for arranging that algebraically cancelling intersection points occur on the same Whitney disc.

Construction 7.1 (Transfer move) *Let Σ and M be as in Convention 1.1 and let $H: (\Sigma, \partial\Sigma) \rightarrow (M, \partial M)$ be a generic immersion, with components $\{h_i: \Sigma_i \rightarrow M\}$. Assume the double points within H are paired by a convenient collection ${}^{\circ}W$ of Whitney discs.*

Let W_1 and W_2 be components of \mathcal{W} with $\text{Int } W_1 \pitchfork H \neq \emptyset \neq \text{Int } W_2 \pitchfork H$. We can perform three finger moves on H , so that the resulting generic immersion H' has six new double points, paired by three framed, embedded Whitney discs $\{V, U_1, U_2\}$, each of which has two intersections with H' , such that the boundaries of $\{V, U_1, U_2\}$ are mutually disjoint and embedded. Moreover, the collection $\mathcal{W}' := \mathcal{W} \cup \{V, U_1, U_2\}$ is a convenient collection of Whitney discs for H' and we have

$$|\text{Int } W_1 \pitchfork H'| = |\text{Int } W_1 \pitchfork H| - 1 \quad \text{and} \quad |\text{Int } W_2 \pitchfork H'| = |\text{Int } W_2 \pitchfork H| - 1.$$

Proof Suppose that W_1 pairs intersections of h_a and h_b while W_2 pairs intersections of h_c and h_d , where repetition within a, b, c and d is allowed. Perform a finger move between h_a and h_c , creating two new double points paired by a corresponding framed, embedded Whitney disc V . Note that the interior of V is disjoint from the image of H . The operation depicted in Figure 8 gives a further regular homotopy, involving a finger move pushing h_e through h_a , and a finger move pushing h_f through h_c . We call the outcome of all three finger moves H' . The procedure creates six new intersections within H' compared with H . The four intersections created by the $h_e - h_a$ and $h_f - h_c$ finger moves are paired by Whitney discs U_1 and U_2 . A preliminary version of these are shown in Figure 8(ii); the final versions are those arising after the boundary push-off operations indicated by the bottom panel. Overall, the move transfers an intersection of H with W_1 , as well as an intersection of H with W_2 , on to V , so that $|\text{Int } V \pitchfork H'| = 2$. By construction, each U_i intersects H' twice. \square

Now we prove Lemma 5.4, whose statement we recall.

Lemma 5.4 *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. Suppose that F admits algebraically dual spheres, and that all double points of F are paired by a convenient collection \mathcal{W} of Whitney discs. Let $\mathcal{W}^\infty \subseteq \mathcal{W}$ denote the subcollection of Whitney discs for the intersections within F^∞ , where F^∞ is as in Definition 5.1. If $t(F^\infty, \mathcal{W}^\infty) = 0$, then $\text{km}(F) = 0$.*

Proof By applying the geometric Casson lemma (Lemma 4.2) and Propositions 2.18, 2.25 and 3.1, we may arrange by a regular homotopy that F and G become geometrically dual. By definition,

$$(7-1) \quad t(F^\infty, \mathcal{W}^\infty) = \sum_{l,i} |\text{Int } W_l^\infty \pitchfork f_i^\infty| = 0 \in \mathbb{Z}/2.$$

We modify the collection of Whitney discs, as follows, so that each has an even number of intersections with F^∞ . Since the count in (7-1) is zero, the number of Whitney discs with odd intersection with F^∞ is even, so we may pair them up (arbitrarily). For each such pair, apply Construction 7.1. This changes F by finger moves to some F' , whose double points are paired by a convenient collection of Whitney discs $\mathcal{W}' := \{W'_i\}$ such that each element of \mathcal{W}' has an even number of intersections with $(F')^\infty$. Note that the new Whitney discs created by the application of Construction 7.1 have been added to the collection.

For each intersection of some W'_i with $(F')^\infty$, tube W'_i into the corresponding geometrically dual sphere. Note that each sphere being tubed into is necessarily twisted, but since we tube an even number of times,




Figure 8: The transfer move. (i) Whitney discs W_1 and W_2 pairing intersections between h_a and h_b , and between h_c and h_d , respectively. (ii) A finger move between h_a and h_c creates a new pair of intersections, paired by a Whitney disc V , shown on both panels. Depicted in the left panel, an intersection between W_1 and h_e is pushed down into h_a and then one of the resulting intersection points is pushed across to V . In the right panel, we see this new intersection between V and h_e . Further, an intersection between W_2 and h_f is pushed down to h_c and one of the resulting intersection points is pushed over to V . These last three moves form a regular homotopy of H , with result H' . Each W_i has one fewer intersection with H' than with H , at the expense of creating two new intersections within H' , paired by Whitney discs U_1 and U_2 , both shaded grey. Additionally, V has two intersections with H' . The result of a boundary push-off operation making the Whitney arcs for the U_i (purple) disjoint from the Whitney arcs for W_i and V is shown in (iii). This operation creates two intersections of each U_i with H' .

the total change in the framing of W'_i is even. Do this for each element of ${}^{\circ}W'$. The resulting family of Whitney discs may still intersect F' , but not $(F')^{\circ}$. For each such intersection with F' , again tube into the appropriate geometrically dual sphere. Now the spheres are not twisted, so the framing of the Whitney discs changes by an even number. Arrange for all the Whitney discs to be framed by adding local cusps in

the interior. We may do this because the framing coefficient of each of the Whitney discs is even. We have now produced the desired convenient collection of Whitney discs for the intersections within F' , whose interiors lie in the complement of F' . This shows that $\text{km}(F') = 0$, and therefore $\text{km}(F) = 0$, as desired. \square

Before giving the proof of Lemma 5.11, we explain the key new construction in this paper, which we already mentioned in Section 1.2. Briefly, given a band B with boundary lying on an immersed surface, a finger move along a fibre of the band produces two new double points paired by a Whitney disc arising from B .

Construction 7.2 (Band fibre finger move) *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. Suppose that $\mu(F) = 0$ and that the self-intersections of F are paired by a convenient collection ${}^{\circ}\mathcal{W} = \{W_I\}$ of Whitney discs with boundary arcs A .*

Consider $B \in \mathfrak{B}(F)$ as a D^1 -bundle. Then we can do a finger move along a fibre of B , with endpoints missing A , as depicted in Figure 1. We call this fibre the **finger arc**, and denote it by D . We assume that D misses all double points of B and all intersections between $\text{Int } B$ and F .

A finger move depends on a choice of 2-dimensional subbundle of the normal bundle to the finger arc (the proof of Lemma 7.3 will give further details). We use a subbundle that lies in the tangent bundle TF at one end of the arc, contains ν_D^B along D , and intersects TF in the line bundle ν_D^B at the other end of the arc. Here ν_D^B is the normal bundle of D in B .

Call the immersion resulting for the above finger move F' . We will check in the next paragraph that the remainder of B , ie the complement in B of a tubular neighbourhood of D , gives a Whitney disc for the new pair of double points. Make the boundary embedded and disjoint from A , by boundary push-off operations, and then boundary twist if necessary, to obtain a framed Whitney disc W_B for the new double points. Then ${}^{\circ}\mathcal{W}' := {}^{\circ}\mathcal{W} \cup \{W_B\}$ is a convenient collection of Whitney discs pairing the double points of F' . Note that both F' and W_B depend on the choice of finger arc D and the 2-dimensional subbundle of its normal bundle mentioned above.

Now, as promised, we check that W_B is a Whitney disc. The finger move creates a trivial Whitney disc and we refer to the corresponding Whitney arcs as the trivial arcs. The double points are also paired by the arcs $A_1, A_2 \subseteq \partial B$, where $A_1 \cup A_2 = \partial W_B$. The existence of the disc W_B implies that the group elements of the double points agree with respect to the arcs A_1 and A_2 . It remains only to see that the double points have opposite signs with respect to the arcs A_1 and A_2 (see Definition 2.27). The case that both M and Σ are orientable is straightforward. The general case follows from the w_1 condition in the definition of a band, (5-1), as we now check.

First we consider the case that B is an annulus. Let $\partial_1 B$ and $\partial_2 B$ denote the two components of ∂B . Then A_1 and A_2 differ from the trivial arcs joining the new double points by $\partial_1 B$ and $\partial_2 B$, respectively. By Definition 2.27, the double points have opposite sign precisely when $\langle w_1(\Sigma), \partial B \rangle + \langle w_1(M), \partial_1 B \rangle = 0$, which matches (5-1) since $\partial_1 B$ is homotopic in M to the core of B .

Now suppose that B is a Möbius band. Then the union of A_1 and A_2 and the trivial pair of arcs is the circle $\partial B \subseteq \Sigma$. Moreover, the union of the image of A_1 and either one of the trivial arcs forms a circle in M homotopic to the core C of B . As before, the double points have opposite sign precisely when $\langle w_1(\Sigma), \partial B \rangle + \langle w_1(M), C \rangle = 0$, which again matches (5-1).

The following lemma explains how Construction 7.2 changes the value of t . In the proof we will also carefully explain how to make suitable choices of 2-dimensional subbundles, as required for the finger move in Construction 7.2.

Lemma 7.3 *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1.*

- (i) *Suppose that F' and \mathcal{W}' are obtained from F and \mathcal{W} by a single application of Construction 7.2 with respect to a band $B \in \mathcal{B}(F)$, where $w_1(\Sigma)$ restricted to every component of ∂B is trivial. Let A denote the Whitney arcs corresponding to \mathcal{W} . Then*

$$t(F', \mathcal{W}') = t(F, \mathcal{W}) + \Theta_A(B) \in \mathbb{Z}/2.$$

- (ii) *Suppose that F' and \mathcal{W}' are obtained from F and \mathcal{W} by a single application of Construction 7.2 with respect to a band $B \in \mathcal{B}(F)$ and an arc $D \subseteq B$, where B is an annulus with $w_1(\Sigma)$ nontrivial on both boundary components and $D \subseteq B$ connects the two boundary components. Let A denote the Whitney arcs corresponding to \mathcal{W} . Then*

$$t(F', \mathcal{W}') = t(F, \mathcal{W}) + \Theta_A(B, D) \in \mathbb{Z}/2.$$

For the proof it will be advantageous to refrain from applying boundary twists and removing intersections involving ∂W_B , and to instead use the following alternative definition of $t(F, \mathcal{W})$, using a slightly weaker restriction on collections of Whitney discs, as in [Stong 1994]; see [Freedman and Quinn 1990, Section 10.8A].

Definition 7.4 *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. A weak collection of Whitney discs for F is a collection \mathcal{W} of Whitney discs pairing all the double points of F , with generically immersed interiors transverse to F , and with Whitney arcs whose interiors are generically immersed in $F(\Sigma)$ minus the double points of F .*

In particular, compared to the definition of a convenient collection of Whitney discs (Definition 2.31), we allow the boundaries of Whitney discs to be generically immersed on Σ and to intersect one another transversely. We also allow the Whitney discs to be twisted, ie for the framing of the normal bundle restricted to the boundary to disagree with the Whitney framing. Each of the discs in a weak collection of Whitney discs admits a normal bundle. The proof is the same as for a generic immersion of pairs, but with a preliminary step that one has to first fix the normal bundle in neighbourhoods of the two double points being paired.

Definition 7.5 Given a weak collection of Whitney discs ${}^{\circ}\mathcal{W} := \{W_l\}$ for the double points of a generic immersion F as in Convention 1.1, fix an ordering on the indexing set for ${}^{\circ}\mathcal{W}$ and define

$$t_{\text{alt}}(F, {}^{\circ}\mathcal{W}) := \sum_l \left(\mu_{\Sigma}(\partial W_l) + \sum_{k>l} |\partial W_l \pitchfork \partial W_k| + \sum_i |\text{Int } W_l \pitchfork f_i| + e(W_l) \right) \in \mathbb{Z}/2,$$

where $e(W_l)$ is the relative Euler number of the normal bundle with respect to the Whitney framing on ∂W_l .

Note that, if ${}^{\circ}\mathcal{W}$ is a convenient collection of Whitney discs for F , then $t_{\text{alt}}(F, {}^{\circ}\mathcal{W}) = t(F, \mathcal{W})$ (see Definition 1.5). In particular, since a convenient collection of Whitney discs comprises framed Whitney discs and has embedded and disjoint Whitney arcs, a majority of terms in the definition of t_{alt} vanish. The following lemma shows that t_{alt} can be used as a proxy for t in general.

Lemma 7.6 *Given a weak collection of Whitney discs ${}^{\circ}\mathcal{W} := \{W_l\}_{l=1}^n$ for the double points of a generic immersion F as in Convention 1.1, there exists a convenient collection of Whitney discs ${}^{\circ}\mathcal{W}'$ such that $t(F, {}^{\circ}\mathcal{W}') = t_{\text{alt}}(F, {}^{\circ}\mathcal{W})$.*

Proof For each Whitney disc W_l with $e(W_l) \neq 0$, add boundary twists to obtain \overline{W}_l with $e(\overline{W}_l) = 0$. Each boundary twist changes the Euler number by ± 1 and introduces an intersection of the Whitney disc with F . We have

$$\sum_i |\text{Int } \overline{W}_l \pitchfork f_i| \equiv \sum_i |\text{Int } W_l \pitchfork f_i| + e(W_l) \pmod{2},$$

and also $\mu_{\Sigma}(\partial \overline{W}_l) = \mu_{\Sigma}(\partial W_l)$ and $|\partial \overline{W}_l \pitchfork \partial \overline{W}_k| = |\partial W_l \pitchfork \partial W_k|$ for each $k \neq l$.

Next we will remove intersections between Whitney arcs as well as self-intersections of Whitney arcs, at the expense of adding intersections between F and the Whitney discs. We will use the procedure described in [Powell and Ray 2021a, Section 15.2.3]. For an intersection between $\partial \overline{W}_l$ and $\partial \overline{W}_k$, where possibly $k = l$, the procedure pushes the intersection off one of the endpoints of one of the Whitney arcs of $\partial \overline{W}_l$, ie a double point of F , moving a neighbourhood of $\partial \overline{W}_k$ and creating an intersection between \overline{W}_k and F . This new intersection point is created in a small neighbourhood of the double point of F chosen for the pushing-off procedure. If several Whitney arcs intersect the given Whitney arc of $\partial \overline{W}_l$, push off in order of proximity to the endpoint. This avoids extraneous intersections between Whitney arcs being created. Perform this pushing-off procedure on both arcs of $\partial \overline{W}_l$. For each of the two arcs in $\partial \overline{W}_l$, push towards one of the two double points of F paired by \overline{W}_l ; choose these double points so that we use one double point for each arc. This ensures that the new intersections between Whitney discs and F arise in disjoint neighbourhoods in the ambient manifold.

Apply the move described in the previous paragraph to the Whitney arcs of $\{\overline{W}_l\}$ in order, beginning with $l = n$. In other words, in the i^{th} step, we push off the intersections of $\partial \overline{W}_k$ with $\partial \overline{W}_{n-i+1}$ for

$k \leq n - i + 1$. After the n^{th} step, we produce a convenient collection $\mathcal{W}' := \{W'_l\}$, where each W'_l is the result of applying the above procedure to \overline{W}_l . This yields, for each l ,

$$\begin{aligned} \sum_i |\text{Int } W'_l \frown f_i| &\equiv \sum_i |\text{Int } \overline{W}_l \frown f_i| + \mu_\Sigma(\partial \overline{W}_l) + \sum_{k>l} |\partial \overline{W}_l \frown \partial \overline{W}_k| \pmod{2} \\ &\equiv \sum_i |\text{Int } W_l \frown f_i| + e(W_l) + \mu_\Sigma(\partial W_l) + \sum_{k>l} |\partial W_l \frown \partial W_k| \pmod{2}. \end{aligned}$$

In the above expression, the term $\sum_{k>l} |\partial \overline{W}_l \frown \partial \overline{W}_k|$ arises since the arcs in $\partial \overline{W}_l$ are moved, to create a new intersection point of \overline{W}_l with F , precisely once for each intersection of $\partial \overline{W}_l$ with $\bigcup_{k>l} \partial \overline{W}_k$.

Sum over l to obtain that $t(F, \mathcal{W}') = t_{\text{alt}}(F, \mathcal{W}) \in \mathbb{Z}/2$, as claimed. □

Proof of Lemma 7.3 By Lemma 7.6, it will suffice to show that, in case (i),

$$t_{\text{alt}}(F', \mathcal{W}') = t_{\text{alt}}(F, \mathcal{W}) + \Theta_A(B) \in \mathbb{Z}/2,$$

and, in case (ii),

$$t_{\text{alt}}(F', \mathcal{W}') = t_{\text{alt}}(F, \mathcal{W}) + \Theta_A(B, D) \in \mathbb{Z}/2.$$

The proof splits into three cases.

Case 1 The band B is an annulus as in (i).

Recall that

$$\Theta_A(B) := \mu_\Sigma(\partial B) + |\partial B \frown A| + |\text{Int } B \frown F| + e(B) \pmod{2}.$$

By the construction of F' and \mathcal{W}' , we have

$$t_{\text{alt}}(F', \mathcal{W}') \equiv t(F, \mathcal{W}) + \mu_\Sigma(\partial W_B) + |\partial W_B \frown A| + |\text{Int } W_B \frown F| + e(W_B) \pmod{2}.$$

Every self-intersection of ∂B and each intersection $\partial B \frown A$ will contribute one intersection of ∂W_B and between ∂W_B and A , respectively, while each intersection $\text{Int } B \frown F$ will contribute one intersection between $\text{Int } W_B$ and F . Thus it remains to show that the framing $e(B)$ appearing in the definition of $\Theta_A(B)$ agrees with the framing $e(W_B)$. For this it will be helpful to pick the finger arc and the 2-dimensional subbundle for its normal bundle needed for the finger move more carefully, which we do next.

Let $\partial_i B$ denote the connected components of $\partial B \subseteq \Sigma$. Consider the following decomposition of the tangent bundle of M restricted to $\partial_i B$:

$$(7-2) \quad TM|_{\partial_i B} \cong T(\partial_i B) \oplus \nu_{\partial_i B}^\Sigma \oplus \nu_{\partial_i B}^B \oplus (\nu_\Sigma^M|_{\partial_i B} \cap \nu_B^M|_{\partial_i B}).$$

As shown in Figure 9, choose a section s of the normal bundle of B that is nonvanishing on the finger arc. In both boundary components $\partial_i B$, we assume that this section lies in $\nu_{\partial_i B}^\Sigma$. Now rotate the section near the top boundary component, as shown in Figure 9, middle, to obtain a section s' , so that on the top component s' lies in $(\nu_\Sigma^M|_{\partial_i B} \cap \nu_B^M|_{\partial_i B})$. For the finger move, by definition, we use the 2-dimensional subbundle of ν_D^M determined by s' and $T(\partial_i B)$, as shown in Figure 9, right.




Figure 9: Left: we see a model annulus B (blue) connecting two sheets of F (black), and a finger arc $D = \{\text{pt}\} \times D^1 \subseteq S^1 \times D^1 \cong B$ (orange). We also see the surface framing on ∂B and the section s along the finger arc of the normal bundle of B in M . Recall the section is defined over all of B , but we only show it on a subset. Middle: in the top half of B , we rotate the section so that it lies in $(\nu_\Sigma^M|_{\partial_i B} \cap \nu_B^M|_{\partial_i B})$, ie in the time direction, on the top boundary component (dotted green). The modified section is called s' . Right: after performing the finger move, s' gives the Whitney framing for the new Whitney disc W_B .

Now consider the Whitney disc W_B obtained from B after performing the finger move along D using the above 2–dimensional subbundle of its normal bundle. By definition, $e(W_B)$ equals the number of zeros of $s'|_{W_B}$, since on the boundary Whitney arcs it is normal to one sheet and tangent to the other. On the other hand, the number of zeros of $s'|_{W_B}$ equals the number of zeros of s , since s' was obtained from s by a rotation, and neither section vanishes near the finger arc. Finally, by definition, $e(B)$ counts the zeros of s . Therefore we see that $e(B) = e(W_B)$.

Case 2 The band B is an annulus such that $w_1(\Sigma)$ restricted to both components of ∂B is nontrivial, as in (ii).

Assume that a finger arc D has been chosen. To define $\Theta_A(B, D)$, we pick a section s as in Definition 5.14 on \widehat{B} , which is by definition the result of cutting B open along D . Recall that

$$\Theta_A(B, D) := \mu_\Sigma(\partial B) + |\partial B \pitchfork A| + |\text{Int } B \pitchfork F| + e(\widehat{B}) \pmod 2.$$

As in Case 1, we need only check that the term $e(\widehat{B})$ in $\Theta_A(B, D)$ agrees with the term $e(W_B)$ in t_{alt} . Again as in Case 1, we rotate the section near the top boundary component to obtain a section s' , so that on the top component s' lies in $(\nu_\Sigma^M|_{\partial_i B} \cap \nu_B^M|_{\partial_i B})$. For the finger move, we use the 2–dimensional subbundle determined by s' and $T(\partial_i B)$. Note that, just like s , the section s' is not defined on all of B , but only on \widehat{B} ; see Figure 10. Nevertheless, on points that map to the same point in D , the section s' agrees up to a sign and thus still determines a 1–dimensional subbundle. The section s' restricts to a section of the normal bundle of the Whitney disc W_B obtained from B . As in Case 1, the quantities $e(B)$ and $e(W_B)$ coincide.

Case 3 The band B is a Möbius band with $w_1(\Sigma)$ restricted to ∂B trivial, as in (i).

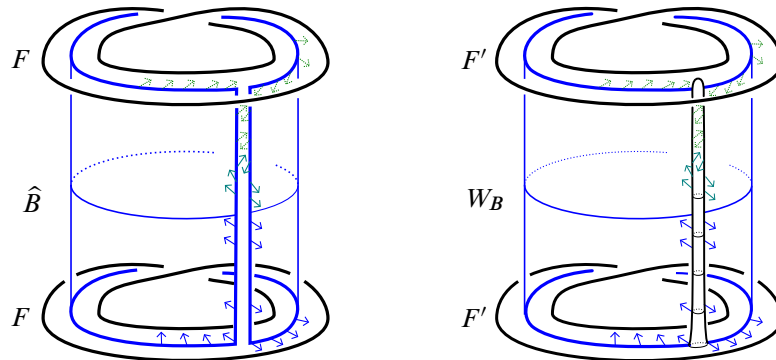


Figure 10: Left: the section s' on D (orange), and on a parallel copy of D ; see Figure 7(b). Close to the top boundary the section extends into the time direction (teal and dotted green). Right: the section s' on the boundary of the new Whitney disc W_B .

As in Case 1, we only need to show that the term $e(B)$ in $\Theta_A(B)$ agrees with the term $e(W_B)$ in t_{alt} . Let D denote a properly embedded arc on B along which we wish to perform the finger move. Identify B with the quotient of the square $S := [-1, 1] \times [-1, 1]$ as usual, ie $(-1, x) \sim (1, -x)$ for all $x \in [-1, 1]$, with D corresponding to the arc $\{-1\} \times [-1, 1] \equiv \{1\} \times [-1, 1]$ (see Figure 11). Pull back the normal bundle ν_B^M of B in M to S via the quotient map $\pi: S \rightarrow B$ and then pick a trivialisation $\pi^*(\nu_B^M) \cong S \times \mathbb{R}^2$ so that, on the horizontal boundary $H := [-1, 1] \times \{-1, 1\}$, we have that $\pi^* \nu_{\partial B}^\Sigma$ coincides with $H \times \mathbb{R} \times \{0\} \subseteq S \times \mathbb{R} \times \mathbb{R}$.

Then we pick a section s of ν_B^M such that $s|_{\partial B}$ lies in $\nu_{\partial B}^\Sigma$. Note that $s|_{\partial B}$ is nowhere-vanishing but the section s might have zeros. Without loss of generality we assume that any zeros of s do not lie in the strip $([-1, -1 + \varepsilon] \cup [1 - \varepsilon, 1]) \times [-1, 1]$ for some $\varepsilon \in (0, \frac{1}{4})$.

Next we modify the section s . Choose a “necklace” region, ie a subsquare N with two opposite edges coinciding with $[-1, -1 + \varepsilon] \times \{1\}$ and $[1 - \varepsilon, 1] \times \{1\}$, and otherwise lying in the interior of $S := [-1, 1]^2$.




Figure 11: (a) The necklace region $N \subseteq S$ is shown in grey. The dotted black lines indicate the boundary of the region where the finger move occurs. The solid black lines indicate where the band is attached to the surface F . The arc D is shown in orange. (b) The section s is shown in blue. (c) The section s' is indicated. Note that on S_- , the sections s and s' agree. On S_+ , the section s' (green) is obtained by rotating s by 90 degrees. In the necklace region, the section rotates continuously (teal), interpolating between the values on S_+ and S_- . Note that the section on the arc D has not changed, but it has been modified on part of the dotted lines.

We consider the pullback of s to S , where we have a trivialisation. Modify this pullback so that it remains unchanged in the lower component S_- of $S \setminus N$ and is rotated by 90 degrees on the upper component S_+ . On the region N , the section rotates continuously, interpolating between the values on S_+ and S_- . Push this forward to get a modified section s' on B . Since the modification is produced by a continuous rotation, the number of zeros of this modification agrees with the number of zeros of the original s .

Recall that we wish to perform a finger move guided by the arc $D = \{-1\} \times [-1, 1]$. Without loss of generality, we assume that the “width” of the finger move is 2ε . More precisely, to perform a finger move we need a 2-dimensional subbundle of ν_D^M . We require that this contains ν_D^B to ensure that the finger move cuts B open into a Whitney disc, as desired. Fix an identification of the total space of the normal bundle ν_D^M with $D \times \mathbb{R}^3$. We choose an embedding $\iota: \nu_D^M \hookrightarrow M$ restricting to the inclusion of D on $D \times \{0\}$, with the following properties:

- (i) We assume the first \mathbb{R}^1 factor of $D \times \mathbb{R}^3$ corresponds to ν_D^B and that ι identifies $D \times \{\pm 1\} \times \{0\} \times \{0\}$ with the arcs $\{-1 + \varepsilon, 1 - \varepsilon\} \times [-1, 1]$. This is what was meant by the width of the finger move.
- (ii) We also require that ι identifies $D \times \{t\} \times \{1\} \times \{0\}$ with $s'(\iota(D \times \{t\} \times \{0\} \times \{0\}))$ for $t \geq 0$, and with $s'(\iota(D \times \{t\} \times \{0\} \times \{0\}))$ rotated by 90 degrees for $t \leq -1$. (Here we also implicitly identify ν_B^M with its image in M .)

Now do the finger move using $D \times S^1 \times \{0\}$ according to this parametrisation, where S^1 is the unit circle in the \mathbb{R}^2 factor, along with a finger tip. The choices above imply that $s'|_{\partial W_B}$ is a Whitney framing, where W_B denotes the new Whitney disc created according to Construction 7.2. Specifically, let F' denote the result of the finger move. By our choice of the 2-dimensional subbundle for the finger move above, the section s' is normal to F' along half of ∂W_B , and tangent along the other half; see Figure 12.

As previously mentioned, we need to check that the relative Euler number $e(B)$ in $\Theta_A(B)$ agrees with the twisting number $e(W_B)$ in t_{alt} . The relative Euler number $e(B)$ is given by the number of zeros of the section s on the interior of B . As mentioned before, this coincides with the number of zeros of the section s' . Since $s'|_{\partial W_B}$ is a Whitney framing, this further coincides with the twisting number $e(W_B)$, as desired, since we assumed there are no zeros of s within the strip $([-1, -1 + \varepsilon] \cup [1 - \varepsilon, 1]) \times [-1, 1] \subseteq B$ used for the finger move. \square

Next we prove Lemma 5.11. Here is the statement for the convenience of the reader.

Lemma 5.11 *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1, with $\mu(F) = 0$. If the $\mathbb{Z}/2$ -valued intersection form λ_Σ on $H_1(\Sigma; \mathbb{Z}/2)$ is nontrivial on $\partial\mathcal{B}(F)$, then we can change F by a regular homotopy to F' such that there are convenient collections of Whitney discs \mathcal{W} and \mathcal{W}' for the double points of F and F' , respectively, such that $t(F, \mathcal{W}) \neq t(F', \mathcal{W}')$.*

Moreover, if F has dual spheres and the $\mathbb{Z}/2$ -valued intersection form λ_{Σ^∞} on $H_1(\Sigma^\infty; \mathbb{Z}/2)$ is nontrivial on $\partial\mathcal{B}(F^\infty)$, then $\text{km}(F) = 0$.




Figure 12: (a) The surface F is shown in black, and the Möbius band B in blue. Note this picture is entirely in \mathbb{R}^3 . The necklace region N is in grey, and splits B into two components S_+ and S_- . The finger arc D is in orange, and the width of the finger move is shown with dotted lines. (b) We show the section s in blue. Note that, while there is a rotation along D , there are no zeros of s in the strip between the dotted lines. (c) The modified section s' . Note the section coincides with s on S_- and has been rotated (green) on S_+ . (d) The section s' on the Whitney disc W_B formed after the band fibre finger move. By the construction of the 2-dimensional subbundle of the normal bundle of D used to guide the finger move, the section s' is tangent to F' along the right edge of the finger (corresponding to the right dotted line in (c), where the finger contains part of a Whitney arc of W_B), and s' is normal to F' along the left edge of the finger.

Proof We first prove the statement (without using dual spheres) about $t(F, \mathcal{W})$ depending on the choice of \mathcal{W} under our assumption. By hypothesis, F is a generic immersion whose double points can be paired by a convenient collection $\mathcal{W} = \{W_i\}$ of Whitney discs (Corollary 2.30). By hypothesis, λ_Σ is nontrivial on $\partial\mathcal{B}(F)$, meaning that there are bands B_1 and B_2 with boundary on $F(\Sigma)$ minus double points such that $\lambda_\Sigma(\partial B_1, \partial B_2) \neq 0 \in \mathbb{Z}/2$. Here it is possible that B_2 is a parallel push-off of B_1 . Using B_i and Construction 7.2, perform a finger move and obtain a new framed Whitney disc, calling the resulting convenient collection of Whitney discs \mathcal{W}_i , for $i = 1, 2$, and the resulting map F_i

If $t(F_i, \mathcal{W}_i) \neq t(F, \mathcal{W})$ for some $i = 1, 2$, we can set $F' = F_i$ and $\mathcal{W}' = \mathcal{W}_i$. Otherwise, use Lemma 7.3 twice, for B_1 and B_2 simultaneously, and let \mathcal{W}' denote the resulting convenient collection of Whitney discs for the resulting map F' . Then the change in $t(F, \mathcal{W})$ is as before, except that there is an additional contribution from the odd number of intersections between the boundary arcs for the new Whitney discs coming from B_1 and B_2 . Specifically, removing these by pushing one Whitney arc off the end of the

other (as part of Construction 7.2) introduces an odd number of intersections between the Whitney discs and F . Therefore, $t(F', \mathcal{W}') \neq t(F, \mathcal{W})$, as needed.

For the second statement, apply the above argument to the subcollection ${}^{\circ}\mathcal{W}^{\circ}$ of \mathcal{W} pairing the intersections within F° . It follows that we may find ${}^{\circ}\mathcal{W}^{\circ}$ such that $t(F^{\circ}, {}^{\circ}\mathcal{W}^{\circ}) = 0$, possibly after a regular homotopy of F° . Then $\text{km}(F) = 0$ follows from Lemma 5.4. \square

For the proof of Lemma 5.16, we will need the following Lemmas 7.7, 7.8, 7.9 and 7.10.

Lemma 7.7 *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. Every element of $H_2(M, \Sigma; \mathbb{Z}/2)$ can be represented by an immersion of some compact surface into M , with interior transverse to F , and with boundary generically immersed in $F(\Sigma)$ away from the double points.*

Proof Let $\mathcal{N}_k(M, \Sigma)$ denote the k -dimensional unoriented bordism group over (M, Σ) , and let \mathcal{N}_k denote the k -dimensional unoriented bordism group over a point. Using topological transversality, it suffices to show that every element of $H_2(M, \Sigma; \mathbb{Z}/2)$ can be represented by a map $(S, \partial S) \rightarrow (M, \Sigma)$ for some surface S . To show this, it suffices to see that the edge homomorphism $\mathcal{N}_2(M, \Sigma) \rightarrow H_2(M, \Sigma; \mathbb{Z}/2)$ from the Atiyah–Hirzebruch spectral sequence is onto.

Recall that the \mathcal{N}_0 is isomorphic to $\mathbb{Z}/2$ while the \mathcal{N}_1 vanishes. It follows that, in the Atiyah–Hirzebruch spectral sequence with E_2 -term $H_p(M, \Sigma; \mathcal{N}_q)$ and converging to $\mathcal{N}_{p+q}(M, \Sigma)$, there is no nontrivial differential going out of $H_2(M, \Sigma; \mathcal{N}_0) \cong H_2(M, \Sigma; \mathbb{Z}/2)$; such a differential would have codomain $H_0(M, \Sigma; \mathcal{N}_1) = 0$. Thus the edge homomorphism $\mathcal{N}_2(M, \Sigma) \rightarrow H_2(M, \Sigma; \mathbb{Z}/2)$ is onto, as desired. \square

Lemma 7.8 *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. Suppose that $\mu(F) = 0$ and let A be a choice of Whitney arcs pairing the double points of F . Then the function Θ_A is quadratic with respect to the $\mathbb{Z}/2$ -valued intersection form λ_{Σ} . That is, let S and S' be compact surfaces, with generic immersions of pairs $(S, \partial S) \looparrowright (M, \Sigma)$ and $(S', \partial S') \looparrowright (M, \Sigma)$ such that ∂S and $\partial S'$ intersect A and each other transversely, and are such that $w_1(\Sigma)$ is trivial on every component of ∂S and $\partial S'$. Then we have*

$$\Theta_A(S \cup S') = \Theta_A(S) + \Theta_A(S') + \lambda_{\Sigma}(\partial S, \partial S').$$

Proof The term $e(S)$ in Definition 5.12 is defined componentwise and the terms $|\partial S \pitchfork A|$ and $|\text{Int } S \pitchfork F|$ are linear in ∂S and S , respectively. Hence, the only term that is not linear in S is $\mu_F(\partial S)$. This term is also quadratic in the sense that

$$\mu_{\Sigma}(\partial S \cup \partial S') = \mu_{\Sigma}(\partial S) + \mu_{\Sigma}(\partial S') + \lambda_{\Sigma}(\partial S, \partial S'),$$

which proves the lemma. \square

Lemma 7.9 *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. Suppose that $\mu(F) = 0$ and let A be a choice of Whitney arcs pairing the double points of F . Let S be a compact surface with a generic immersion of pairs $(S, \partial S) \looparrowright (M, \Sigma)$ such that ∂S is transverse to A and $w_1(\Sigma)$ is trivial*




Figure 13: Adding a disc D to S to remove a self-intersection of ∂S . Left: the neighbourhood of a self-intersection of ∂S before adding the disc D . The section γ_S is shown along ∂S . Middle: the boundary of the disc D . The section γ_D is shown along ∂D . Right: after the modification; see Figure 14.

on every component of ∂S . Then there is another such surface S' with a generic immersion of pairs $(S', \partial S') \looparrowright (M, \Sigma)$ with $\partial S'$ transverse to A such that

- (1) $[S] = [S'] \in H_2(M, \Sigma; \mathbb{Z}/2)$;
- (2) $\Theta_A(S) = \Theta_A(S')$; and
- (3) $\partial S'$ is embedded in Σ .

Proof To start, pick a section γ_S of the normal bundle ν_S^M which on ∂S is nowhere-vanishing and lies in $\nu_{\partial S}^F$ as in the definition of $\Theta_A(S)$.

The idea of the proof is to remove all intersections of ∂S by locally adding a twisted disc D as indicated in Figure 13. More precisely, we add these discs D such that the interiors are disjoint from the interior of F and the boundary is disjoint from A . Then pick a section γ_D of ν_D^M such that, along the aligned (ie parallel) parts of the boundaries, γ_D and γ_S are opposite. Glue D to S along the aligned parts of the boundaries and push this part of the boundary off F as indicated in Figure 14. Each of these local twisted discs has mod 2 Euler number 1, as can be seen from the nontrivial linking in Figure 15. Thus the resulting surface S' has embedded boundary and the mod 2 Euler number of S' differs from that of S by the number of intersections of ∂S modulo two, ie $\mu_\Sigma(\partial S)$. Since we have changed neither the




Figure 14: Glue D to S along the aligned parts of the boundaries and push this part of the boundary off F .




Figure 15: A twisted band with Euler number +1 in a movie description. Bottom: an immersed figure-eight curve (blue) is shown lying on the immersed surface $F(\Sigma)$ (black) away from the double points. A framing on the normal bundle on the boundary of the band is shown in light blue. Moving upward/forward in time, we see a simple closed curve shrinking to a point. The push-off corresponding to the framing induced by Σ is shown in light blue. For the twisted band with Euler number -1 , we use the other resolution.

number of intersections of the interior with F nor the number of intersections of the boundary with A , we have $\Theta_A(S) = \Theta_A(S') \in \mathbb{Z}/2$. As the local discs are trivial in $H_2(M, \Sigma; \mathbb{Z}/2)$, we furthermore have $[S] = [S'] \in H_2(M, \Sigma; \mathbb{Z}/2)$, as needed. \square

Lemma 7.10 *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. Let Z be a disjoint union of embedded circles in Σ . Let $\Sigma | Z$ denote Σ cut along Z , ie the completion of $\Sigma \setminus Z$ to a compact manifold with boundary. Let $\mathcal{F} = \{\Sigma_i\}$ be the connected components of $\Sigma | Z$ and suppose that $[Z] = 0 \in H_1(\Sigma; \mathbb{Z}/2)$. Then we can pick a subset $\mathcal{F}' \subseteq \mathcal{F}$ such that each component of Z appears exactly once as a connected component of the boundary of precisely one $\Sigma_i \in \mathcal{F}'$.*

Proof Without loss of generality, assume that Σ is connected. Considering the entire collection $\mathcal{F} = \{\Sigma_i\}$, every component of Z would appear as the boundary of precisely two of the Σ_i , since otherwise Z would be nontrivial in $H_1(\Sigma; \mathbb{Z}/2)$. To see this, note that Z can contain homologically essential curves in $H_1(\Sigma; \mathbb{Z}/2)$ provided they cancel. However, none of these can be orientation-reversing curves, since Z is embedded.

The idea of the proof is to take “half” of the components of \mathcal{F} . Let $x \in \Sigma | Z$ be an arbitrary basepoint away from Z . For each Σ_i , define $p(\Sigma_i) \in \mathbb{Z}/2$ as follows. Pick a point $y \in \text{Int } \Sigma_i$ and a path w in Σ from x to y which is transverse to Z . Define $p(\Sigma_i)$ as the mod 2 intersection number of w and Z .

We show that $p(\Sigma_i)$ is independent of the choices of w and y . If w' is another path from x to y , then the concatenation $w^{-1} \cdot w'$ is a loop in Σ and we have

$$|(w^{-1} \cdot w') \cap Z| = \lambda_\Sigma([w^{-1} \cdot w'], [Z]) = \lambda_\Sigma([w^{-1} \cdot w'], 0) = 0.$$

So $p(\Sigma_i)$ does not depend on the choice of w . Also, since each Σ_i is connected, $p(\Sigma_i)$ does not depend on y . To see this, let $y' \in \text{Int } \Sigma_i$ and choose a path z from y to y' that lies in $\text{Int } \Sigma_i$. Let w' be a path

from x to y' which is further transverse to Z . Then

$$|w \pitchfork Z| = |(w \cdot z) \pitchfork Z| = |w' \pitchfork Z|.$$

The first equation uses that $z \subseteq \Sigma_i$ and the second uses independence of the choice of w . Hence, $p(\Sigma_i) \in \mathbb{Z}/2$ is well defined, as desired.

Now let \mathcal{F}' consist of all the components Σ_i for which $p(\Sigma_i) = 0$. This subset is the one we seek, since, for a fixed component Z_j of Z , the two components of \mathcal{F} containing a cut-open copy of Z_j have different values of p . □

We are now ready for the proof of Lemma 5.16.

Lemma 5.16 *Let $F : (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1, with $\mu(F) = 0$. Let A be a choice of Whitney arcs pairing the double points of F .*

- (i) *Let S be a compact surface with a generic immersion of pairs $(S, \partial S) \looparrowright (M, \Sigma)$ such that ∂S is transverse to A and $w_1(\Sigma)$ is trivial on every component of ∂S . Then $\Theta_A(S) \in \mathbb{Z}/2$ depends only on the homology class of S in $H_2(M, \Sigma; \mathbb{Z}/2)$.*
- (ii) *Let B be an annulus with a generic immersion of pairs $(B, \partial B) \looparrowright (M, \Sigma)$ such that ∂B is transverse to A and $w_1(\Sigma)$ is nontrivial on both components of ∂B . Pick an embedded arc D in B connecting the components of ∂B and disjoint from all double points. Then $\Theta_A(B, D) \in \mathbb{Z}/2$ depends only on the homology class of B in $H_2(M, \Sigma; \mathbb{Z}/2)$. In particular, $\Theta_A(B, D)$ does not depend on D , so we write $\Theta_A(B)$.*
- (iii) *Let S be a surface as in (i) and let B be an annulus as in (ii) such that $[S] = [B] \in H_2(M, \Sigma; \mathbb{Z}/2)$. Then $\Theta_A(S) = \Theta_A(B) \in \mathbb{Z}/2$.*
- (iv) *If $\lambda_\Sigma|_{\partial\mathcal{B}(F)} = 0$, the restriction of Θ_A to $\mathcal{B}(F)$ is independent of the choice of A , giving a well-defined map $\Theta : \mathcal{B}(F) \rightarrow \mathbb{Z}/2$.*

Proof To prove (i), assume that S and S' are immersed compact surfaces, with $w_1(\Sigma)$ trivial on each of the connected components of the boundaries, representing the same element in $H_2(M, \Sigma; \mathbb{Z}/2)$. Modulo isotopy we can assume that S and S' intersect transversely in their interiors in M , and their boundaries intersect transversely on F . In particular, their boundaries ∂S and $\partial S'$ intersect in an even number of points. Hence, $\Theta_A(S \cup S') = \Theta_A(S) + \Theta_A(S')$ by Lemma 7.8. Thus it suffices to show that $\Theta_A(S) = 0$ for a compact surface S such that $0 = [S] \in H_2(M, \Sigma; \mathbb{Z}/2)$ and $w_1(\Sigma)$ is trivial on ∂S . In particular, we know by Lemma 7.7 that every element of $H_2(M, \Sigma; \mathbb{Z}/2)$ — in particular the trivial class — can be represented by an immersed surface S .

By Lemma 7.9, we can assume that ∂S is embedded. As $0 = [S] \in H_2(M, \Sigma; \mathbb{Z}/2)$, we also have that $0 = [\partial S] \in H_1(\Sigma; \mathbb{Z}/2)$ since S maps to ∂S under the map $H_2(M, \Sigma; \mathbb{Z}/2) \rightarrow H_1(\Sigma; \mathbb{Z}/2)$. Pick a set \mathcal{F}' of components of $\Sigma \setminus \partial S$ as in Lemma 7.10. Gluing the $F_i \in \mathcal{F}'$ to S along the common boundary, we obtain a closed surface N . First note that N represents the same class as S in $H_2(M, \Sigma; \mathbb{Z}/2)$ since

it only differs by a subset of $F(\Sigma)$. Hence, $0 = [N] \in H_2(M, \Sigma; \mathbb{Z}/2)$. As N is closed, it also defines an element in $H_2(M; \mathbb{Z}/2)$. Note that we have the pair sequence

$$\dots \rightarrow H_2(\Sigma; \mathbb{Z}/2) \xrightarrow{F} H_2(M; \mathbb{Z}/2) \rightarrow H_2(M, \Sigma; \mathbb{Z}/2) \rightarrow \dots$$

Hence, N represents the same class in $H_2(M; \mathbb{Z}/2)$ as a subsurface Σ' of Σ . Let λ_M denote the $\mathbb{Z}/2$ -valued intersection form on $H_2(M; \mathbb{Z}/2)$. By hypothesis, we have $\lambda_M(f_j, f_{j'}) = 0$ for any two connected components f_j and $f_{j'}$ of F . Thus,

$$(7-3) \quad \lambda_M([N], [F]) + \lambda_M([N], [N]) = 0.$$

We finish the proof of (i) by showing that $\Theta_A(S) = \lambda_M([N], [F]) + \lambda_M([N], [N])$.

Recall that we were able to assume that ∂S is embedded in $F(\Sigma)$ away from the double points and that $\lambda_M([N], [N]) = e(\nu N) \pmod 2$. We claim that this in turn agrees with $e(S) + \sum_{F_i \in \mathcal{F}'} e(F_i)$. Here we define $e(F_i)$ as follows. We used F to define a nowhere-vanishing section of $\nu_S^M|_{\partial S}$. Since $\nu_S^M|_{\partial S}$ is 2-dimensional, we can pick a linearly independent nonvanishing section. This can be equivalently used for the definition of $e(S)$. But this new section now can also be used to define $e(F_i)$. Combining these vector fields that are transverse to the zero section defines a vector field on the normal bundle of N , and hence computes the Euler number of the normal bundle of N . Thus we have shown

$$\lambda_M([N], [N]) = e(S) + \sum_{F_i \in \mathcal{F}'} e(F_i).$$

Now consider $\lambda_M([N], [F])$. We can use the vector field used for defining $e(F_i)$ to make N and F transverse. Then $\lambda_M([N], [F])$ is given by the sum of $|S \pitchfork F|$, $\sum_{F_i \in \mathcal{F}'} e(F_i)$ and the self-intersection points of F contained in the $F_i \in \mathcal{F}'$. As the self-intersection points of F are paired by the Whitney arcs A , we have that modulo two the number of self-intersection points of F contained inside F_i agrees with $|A \pitchfork \partial F_i|$. Since the boundary of the F_i is precisely ∂S , we have

$$\lambda_M([N], [F]) = |\text{Int } S \pitchfork F| + \sum_{F_i \in \mathcal{F}'} e(F_i) + |A \pitchfork \partial S|.$$

Therefore,

$$\begin{aligned} \lambda_M([N], [N]) + \lambda_M([N], [F]) &= e(S) + \sum_{F_i \in \mathcal{F}'} e(F_i) + |\text{Int } S \pitchfork F| + \sum_{F_i \in \mathcal{F}'} e(F_i) + |A \pitchfork \partial S| \\ &= e(S) + |\text{Int } S \pitchfork F| + |A \pitchfork \partial S| \\ &= \Theta_A(S) \in \mathbb{Z}/2, \end{aligned}$$

where the last equality holds because $\mu_S(\partial S) = 0$. Combine this with (7-3) to obtain

$$\Theta_A(S) = \lambda_M([N], [N]) + \lambda_M([N], [F]) = 0.$$

This completes the proof of (i).




Figure 16: A strip, ie half of the tube, added to Δ .

Before proving (ii) and (iii), we introduce a general construction. Let B be an annulus as in (ii) with an embedded arc D in B connecting its two boundary components. As in Remark 5.15, add a tube to $F(\Sigma)$ along the arc D . Let d, d' denote the two discs removed from F when the tube is added. Adding the tube changes F to some F' , an immersion of a surface Σ' , and changes B to a disc Δ . As before, observe that $\partial\Delta$ is an orientation-preserving curve in Σ' and A is now a collection of Whitney arcs pairing the double points of F' . By construction, we see that $\Theta_A(B, D) = \Theta_A(\Delta)$.

Moreover, suppose there is either some immersed compact surface S in M as in (i) or some immersed annulus B' as in (ii), with an embedded arc D' on B' connecting its two boundary components, where possibly $B = B'$. We may choose the tube in the above construction thin enough that $\Theta_A(S)$ and $\Theta_A(B', D')$ remain unchanged. In particular, this means we assume, after a small local isotopy, that the discs d and d' do not intersect the boundaries of S and B' , so both represent classes in $H_2(M, \Sigma'; \mathbb{Z}/2)$. We have the following claim.

Claim 7.11 *If $[B] = [S] \in H_2(M, \Sigma; \mathbb{Z}/2)$, then either $[\Delta] = [S] \in H_2(M, \Sigma'; \mathbb{Z}/2)$ or $[\Delta] = [S] + [d] \in H_2(M, \Sigma'; \mathbb{Z}/2)$. Similarly, if $[B] = [B'] \in H_2(M, \Sigma; \mathbb{Z}/2)$ then either $[\Delta] = [B'] \in H_2(M, \Sigma'; \mathbb{Z}/2)$ or $[\Delta] = [B'] + [d] \in H_2(M, \Sigma'; \mathbb{Z}/2)$.*

Proof The exact sequence of the triple with $\mathbb{Z}/2$ coefficients yields

$$(\mathbb{Z}/2)^2 \cong H_2(\Sigma, \Sigma \setminus (\overset{\circ}{d} \cup \overset{\circ}{d}')) \rightarrow H_2(M, \Sigma \setminus (\overset{\circ}{d} \cup \overset{\circ}{d}')) \xrightarrow{j} H_2(M, \Sigma) \rightarrow H_1(\Sigma, \Sigma \setminus (\overset{\circ}{d} \cup \overset{\circ}{d}')) = 0,$$

so j is surjective with kernel generated by the images of $[d]$ and $[d']$ from the left-hand group.

Construct a lift \tilde{B} of Δ in $H_2(M, \Sigma \setminus (\overset{\circ}{d} \cup \overset{\circ}{d}'))$ by adding a strip along the added tube to Δ , as shown in Figure 16. Since \tilde{B}, S and B' are mapped by j to B, S and B' in $H_2(M, \Sigma)$, respectively, and the kernel is generated by $[d]$ and $[d']$, we see that the classes of \tilde{B}, S and B' differ at most by the classes $[d]$ and $[d']$. The map $H_2(M, \Sigma \setminus (\overset{\circ}{d} \cup \overset{\circ}{d}')) \rightarrow H_2(M, \Sigma')$ identifies $[d]$ and $[d']$, so the claim follows. \square

We continue now to prove (ii). Let B and B' be immersed annuli in M as in the statement of (ii). Choose arcs D in B and D' in B' connecting the boundary components of each, and assume that $[B] = [B']$. By the construction from the proof of Claim 7.11 applied twice, once to B and once to B' , we find discs Δ and Δ' , coming from B and B' , respectively, such that $\Theta_A(B, D) = \Theta_A(\Delta)$ and $\Theta_A(B', D') = \Theta_A(\Delta')$.

From Claim 7.11, applied twice with the roles of B and B' reversed, we see that the classes $[\Delta]$ and $[\Delta']$ satisfy

$$[\Delta] = [B] \quad \text{or} \quad [\Delta] = [B] + [d],$$

and

$$[\Delta'] = [B'] \quad \text{or} \quad [\Delta'] = [B'] + [d].$$

Since also $[B] = [B']$, it follows that either $[\Delta] = [\Delta']$ or $[\Delta] = [\Delta'] + [d]$ in $H_2(M, \Sigma'; \mathbb{Z}/2)$ for Σ' the surface obtained from applying the construction (twice) to Σ .

In the first case, $[\Delta] = [\Delta']$, the proof of (ii) is completed by appealing to (i), which says that $\Theta_A(\Delta) = \Theta_A(\Delta')$, since both Δ and Δ' have $w_1(\Sigma')$ trivial on the boundary. In the second case, $[\Delta] = [\Delta'] + [d]$, we also appeal to (i), but now for the pair of surfaces Δ and $\Delta' \cup d$. So we have that $\Theta_A(\Delta) = \Theta_A(\Delta' \cup d)$. It follows directly from the definition that $\Theta_A(\Delta' \cup d) = \Theta_A(\Delta')$. This completes the proof of (ii). In particular, we have proved that $\Theta_A(B, D)$ does not depend on the choice of arc D .

The proof of (iii) is similar. Suppose we have an immersed annulus B in M as in the statement of (ii), as well as an immersed compact surface S in M as in the statement of (i). Choose an embedded arc $D \subseteq B$ connecting the boundary components. Assume that $[S] = [B] \in H_2(M, \Sigma; \mathbb{Z}/2)$. Apply the previous construction to B , yielding a disc Δ which, by Claim 7.11, satisfies either $[\Delta] = [S]$ or $[\Delta] = [S] + [d]$ in the group $H_2(M, \Sigma'; \mathbb{Z}/2)$ for the surface Σ' obtained from applying the construction to Σ . Further, we know that $\Theta_A(B, D) = \Theta_A(\Delta)$. Now, in the first case, the proof is completed by appealing to (ii), which says that $\Theta_A(\Delta) = \Theta_A(S)$. In the second case, apply (ii) to the pair Δ and $S \cup d$, to see that $\Theta_A(\Delta) = \Theta_A(S \cup d)$. It follows directly from the definition that $\Theta_A(S \cup d) = \Theta_A(S)$.

It remains to prove (iv). Let B denote an element of $\mathcal{B}(F)$. First note that only the term $|\partial B \pitchfork A|$ of $\Theta_A(B)$ depends on the Whitney arcs A . Let A' denote another collection of Whitney arcs. The quantities $\Theta_A(B)$ and $\Theta_{A'}(B)$ differ by $|\partial B \pitchfork A| + |\partial B \pitchfork A'|$, regardless of whether Σ is orientable or nonorientable.

Case 1 The collections of Whitney arcs A and A' correspond to the same choice of pairing up of the double points of F .

For each pair of double points, we can pick Whitney discs W_1 and W_2 with boundary in A and A' , respectively. By adding small strips to the union of W_1 and W_2 in the neighbourhood of the double points, we can see that the difference of A and A' is the boundary of some collection of bands B' . For more details about this construction, see the upcoming proof of Theorem 1.6. Then we have

$$|\partial B \pitchfork A| + |\partial B \pitchfork A'| = |\partial B \pitchfork (A \cup A')| = |\partial B \pitchfork \partial B'| = \lambda_\Sigma(\partial B, \partial B') \pmod{2},$$

which vanishes by assumption.

Case 2 The collections of Whitney arcs A and A' correspond to a different pairing up of the double points of F .




Figure 17: Left: the Whitney discs W_1 , W_2 and V_1 , pairing up double points as (p_1, p_2) , (q_1, p_2) and (q_1, q_2) , respectively. Right: the Whitney disc V_2 pairing up (p_1, q_2) is obtained as a union of W_1 , W_2 and V_1 , by adding small bands at the points p_2 and q_1 to resolve the singularities, and pushing the interiors of the bands into the complement of F . Compare with [Stong 1994, Figure 2].

From A we can construct Whitney arcs A'' so that A' and A'' correspond to the same pairing up of double points, as in Figure 17. Here are the details. We will define the family A'' iteratively, starting with A . Let p_1, p_2, q_1 and q_2 be double points of F . Suppose that arcs in A pair up p_1 and p_2 , as well as q_1 and q_2 , while arcs in A' pair up p_2 and q_1 . Pick Whitney discs W_1 and W_2 with boundary in A . Let V_1 be a Whitney disc for the points p_2 and q_1 with boundary away from A . Then, as indicated in Figure 17, we may choose Whitney arcs, away from the other arcs in A , so that p_2 and q_1 are also paired by a Whitney disc V_2 , obtained as a union of W_1, W_2 and V_1 . Modify the family A by removing ∂W_1 and ∂W_2 , and adding in ∂V_1 and ∂V_2 . Comparing this new family with A' , we see that we have reduced the number of mismatches in the pairing up of double points of F . Iterate this process and call the result A'' .

Looking more closely at the construction in the previous paragraph, observe that, at each step, the family of arcs changes by adding in two parallel copies of the boundary of a Whitney disc V_1 . Since intersection points are counted modulo 2, Θ_A and $\Theta_{A''}$ are equal. By Case 1, we know that $\Theta_{A'}$ and $\Theta_{A''}$ are equal when restricted to $\mathcal{B}(F)$. Thus, Θ_A and $\Theta_{A'}$ are equal when restricted to $\mathcal{B}(F)$, as needed. \square

8 Proof of Theorems 1.6 and 1.9

First we prove Theorem 1.9 from the introduction, which shows that, for b -characteristic surfaces, $t(F, \mathcal{W}) \in \mathbb{Z}/2$ is well defined, ie independent of the Whitney discs \mathcal{W} . Note that the theorem has no assumption about the existence of algebraically dual spheres.

Theorem 1.9 *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1 with $\mu(F) = 0$. Let \mathcal{W} be a convenient collection of Whitney discs for the double points of F . Then F is b -characteristic if and only if, for every F' regularly homotopic to F and convenient collection \mathcal{W}' for the double points of F' , we have $t(F, \mathcal{W}) = t(F', \mathcal{W}')$.*

For b -characteristic F , we denote the resulting regular homotopy invariant by $t(F) \in \mathbb{Z}/2$. Then, if $\text{km}(F) = 0$ — eg if F is an embedding — then $t(F) = 0$.

Proof The final sentence, that $\text{km}(F) = 0$ implies $t(F) = 0$ for b -characteristic F , is an immediate consequence of the definitions.




Figure 18: Within Σ we see the preimages p_1^\pm and p_2^\pm , for the double points p_1 and p_2 of F , respectively. Blue denotes the Whitney arcs for W_1 while red denotes the Whitney arcs for the new disc V_1 . On the left, the choice of sheets stays the same, while it changes on the right. Compare with [Stong 1994, Figure 3].

Now suppose that F is not b -characteristic. Then, by Lemma 5.11, we can assume that $\lambda_\Sigma|_{\partial\mathcal{B}(F)}$ is trivial, which implies that the function Θ is well defined on $\mathcal{B}(F)$. Since we assume that F is not b -characteristic, there exists $B \in \mathcal{B}(F)$ such that $\Theta(B) = 1$, so we can apply Construction 7.2 and Lemma 7.3 to find F' regularly homotopic to F , and a convenient collection of Whitney discs \mathcal{W}' for the double points of F' with $t(F, \mathcal{W}) \neq t(F', \mathcal{W}')$.

If F is b -characteristic, by definition $\lambda_\Sigma|_{\partial\mathcal{B}(F)}$ is trivial and Θ is trivial on $\mathcal{B}(F)$. As indicated above, the function Θ , as well as which classes of $H_2(M, \Sigma; \mathbb{Z}/2)$ can be represented by bands, only depends on the immersion F up to regular homotopy. We need to show that $t(F, \mathcal{W})$ does not depend on the choice of pairing of the double points, the choice of Whitney arcs, nor the choice of Whitney discs; see Figure 18. Let \mathcal{W} be a given initial choice of convenient collection of Whitney discs for the double points of F . Let A be the corresponding collection of Whitney arcs for the double points of F .

The remainder of the proof is similar to [Stong 1994, pages 1311–1313; Freedman and Quinn 1990, Section 10.8A]. We will work with *weak collections of framed Whitney discs* and the alternative count $t_{\text{alt}} \in \mathbb{Z}/2$, as in Definitions 7.4 and 7.5. So the boundaries of our collections of Whitney discs might not be disjointly embedded, but the Whitney discs will be framed (as can always be arranged by boundary twisting). We will show that $t_{\text{alt}}(F, \mathcal{W})$ does not depend on the choice of weak collection of Whitney discs \mathcal{W} , and then use that $t_{\text{alt}}(F, \mathcal{W}) = t(F, \mathcal{W})$ for \mathcal{W} a convenient collection (Lemma 7.6).

Claim 8.1 *Suppose we are given a weak collection of Whitney discs \mathcal{W} corresponding to some choice of pairing up of double points of F ; then, for any other choice of pairing, there exists a weak collection of Whitney discs \mathcal{V} for that choice such that $t_{\text{alt}}(\mathcal{V}) = t_{\text{alt}}(\mathcal{W})$.*

Proof Let p_1, p_2, q_1 and q_2 be double points of F . Suppose that, in the initial choice of data, p_1 and p_2 are paired by a Whitney disc $W_1 \in \mathcal{W}$, and q_1 and q_2 by a Whitney disc $W_2 \in \mathcal{W}$. Suppose we instead pair up p_1 and q_2 by some Whitney disc V_1 . Then, as indicated in Figure 17, p_2 and q_1 are also paired by




Figure 19: Left: two sheets of the surface Σ and two Whitney discs W_i and V_i between the same pair of double points. The disc W_i is assumed to be framed, embedded, have interior disjoint from F and the Whitney discs $\mathcal{U}_{i-1} \setminus \{W_i\}$, and ∂W_i disjoint from A_i . One of its Whitney arcs a_i is also labelled. The blue strip to the right of a_i is an extension of W_i beyond its boundary, which is part of the data for the Whitney move. Right: the result of the Whitney move. The strip and the disc V_i from the previous panel have formed a band B (blue).

a Whitney disc V_2 , obtained as a union of W_1 , W_2 and V_1 . Then $(\mathcal{W} \setminus \{W_1, W_2\}) \cup \{V_1, V_2\}$ is a weak collection of framed Whitney discs. The contribution of V_1 and V_2 to $t_{\text{alt}}(F, (\mathcal{W} \setminus \{W_1, W_2\}) \cup \{V_1, V_2\})$ counts the intersections of F with each disc W_1 and W_2 once, while it counts the intersections of F with the disc V_1 twice. Each intersection of ∂V_1 with $A \setminus (\partial W_1 \cup \partial W_2)$ can be paired with an intersection of ∂V_2 with $A \setminus (\partial W_1 \cup \partial W_2)$. Each intersection of ∂V_1 with $\partial W_1 \cup \partial W_2$ gives rise to two contributions to t_{alt} : an intersection of ∂V_2 with ∂V_1 and a self-intersection of ∂V_2 . Since intersections are counted mod 2 in the definition of t_{alt} , we see that

$$t_{\text{alt}}(F, (\mathcal{W} \setminus \{W_1, W_2\}) \cup \{V_1, V_2\}) = t_{\text{alt}}(F, \mathcal{W}) \in \mathbb{Z}/2,$$

as needed. Iterate this process to complete the proof of Claim 8.1. □

Continuing with the proof of Theorem 1.9, next we check that t_{alt} is independent of the choice of Whitney discs. This includes potentially changing the Whitney arcs and the choice of sheets at each double point. Suppose we are given another weak collection of framed Whitney discs \mathcal{V} for the double points of F . By applying Claim 8.1, we may assume that \mathcal{V} corresponds to the same pairing of double points of F as \mathcal{W} . Assume the collections are indexed so that $W_i \in \mathcal{W}$ and $V_i \in \mathcal{V}$ correspond to the same pair of double points. For each i , define the weak collection of Whitney discs

$$\mathcal{U}_i := \{V_1, V_2, \dots, V_i, W_{i+1}, W_{i+2}, \dots, W_N\},$$

where $\mathcal{U}_0 = \mathcal{W}$ and $\mathcal{U}_N = \mathcal{V}$. We will show that $t_{\text{alt}}(F, \mathcal{U}_{i-1}) = t_{\text{alt}}(F, \mathcal{U}_i) \in \mathbb{Z}/2$ for each i . Let A_i denote the collection of Whitney arcs for \mathcal{U}_i . First we prove a special case.

Claim 8.2 *Suppose the Whitney disc W_i is framed, embedded, with interior disjoint from F and the Whitney discs $\mathcal{U}_{i-1} \setminus \{W_i\}$, and with ∂W_i disjoint from A_i , other than the endpoints. Then $t_{\text{alt}}(F, \mathcal{U}_{i-1}) = t_{\text{alt}}(F, \mathcal{U}_i) \in \mathbb{Z}/2$.*

Proof A neighbourhood of W_i is depicted in Figure 19. Note that the two arcs of ∂V_i lie in A_i and thus, by hypothesis, only intersect the arcs in ∂W_i at the endpoints. As described in the figure, we wish to

perform the Whitney move using W_i pushing towards the Whitney arc a_i for W_i . Observe that the union of V_i with a strip, corresponding to the unit outward-pointing normal vector field of $a_i \subseteq \partial W_i$, is either an annulus or a Möbius band; this requires a small isotopy of V_i to ensure that the chosen vector field of a_i is compatible with the Whitney arcs of V_i , as shown in Figure 19. Denote the union of V_i and the strip by B .

We show that $B \in \mathcal{B}(F)$. For this we need to check that condition (5-1) holds. From Figure 19, right, one sees that ∂B is homotopic in Σ to the union of ∂V_i and ∂W_i . The core C of B is given by the union of a_i and either of the Whitney arcs of V_i . The Whitney arcs must induce opposite signs at the two double points, as explained in Definition 2.27. The orientation conditions in the latter definition imply that the condition in (5-1) holds, as we explain next. Let p_1 and p_2 denote the double points paired by W_i (and V_i). Let a_i and b_i denote the Whitney arcs of W_i , and let c_i and d_i denote those of V_i . Begin by fixing local orientations of M and both sheets of Σ at p_1 , so that the first agrees with the one determined by the latter two. Transport the local orientations of Σ to p_2 via the Whitney arcs of W_i and form the induced local orientation of M at p_2 . By Definition 2.27, this does not agree with the local orientation of M at p_2 determined by the one at p_1 by transporting along a_i . Continuing with the local orientations at p_2 determined in the previous step, transport the local orientations of Σ back to p_1 , this time along the Whitney arcs of V_i . Again by Definition 2.27, the resulting induced local orientation of M at p_1 agrees with the local orientation of M transported to p_1 along c_i . In this circuit, we have constructed a new set of local orientations of M and the two sheets of Σ at p_1 . Compared to the initial choice, the local orientation induced by the sheets of Σ has changed by $\langle w_1(\Sigma), a_i \cup b_i \cup c_i \cup d_i \rangle = \langle w_1(\Sigma), \partial V_i \cup \partial W_i \rangle$. On the other hand, the local orientation of M transported along $a_i \cup c_i$ has changed by $\langle w_1(M), a_i \cup c_i \rangle = \langle w_1(M), C \rangle$, where C is the core of B from above. Since the two orientations must agree, we have $\langle w_1(\Sigma), \partial V_i \cup \partial W_i \rangle = \langle w_1(M), C \rangle$, as needed.

For the band B as above, performing a finger move as in Construction 7.2 creates W_i as the standard Whitney disc, and V_i as the new Whitney disc arising from the band. Here we used the fact that ∂W_i and ∂V_i only intersect at the endpoints. Since F is b -characteristic, the disc V_i has trivial contribution to $t_{\text{alt}}(F, \mathcal{U}_i)$ by Lemma 7.3. So does W_i to $t_{\text{alt}}(F, \mathcal{U}_{i-1})$ since, by hypothesis, ∂W_i is framed, embedded, and disjoint from $A_{i-1} \setminus \{a_i, b_i\} \subseteq A_i$, and the interior of W_i is disjoint from F . Therefore, $t_{\text{alt}}(F, \mathcal{U}_{i-1}) = t_{\text{alt}}(F, \mathcal{U}_i) \in \mathbb{Z}/2$, as asserted. \square

Now we prove the general case. Denote the double points paired by W_i by p_1 and p_2 . By a small isotopy, assume that, other than p_1 and p_2 , the arcs of ∂W_i intersect the arcs in A_i in isolated double points in the interiors. By performing a suitable finger move near p_2 , split W_i into new Whitney discs W'_i and U_1 , creating two new double points q_1 and q_2 in the process, paired by a standard trivial Whitney disc U_2 , where U_1 satisfies the conditions of Claim 8.2. We choose both the base and tip of the finger arc to be closer to p_2 than any intersections of ∂W_i with arcs in A_i , as well as any self-intersections of ∂W_i . See Figure 20. By construction, the points p_1 and q_1 are paired by W'_i , and the points q_2 and p_2 are paired by U_1 . Here U_1 is framed, embedded, has interior disjoint from F and the Whitney discs




Figure 20: Splitting a Whitney disc W_i into two Whitney discs. One of the new Whitney discs, U_1 , pairing p_2 and q_2 , satisfies the hypotheses of Claim 8.2. The other Whitney disc W'_i intersects whatever W_i intersected. The trivial Whitney disc U_2 pairing the new double points q_1 and q_2 is shown in grey. Note that ∂W_i may intersect ∂V_i , or more generally other arcs in A_i , or itself.

$\mathcal{U}_{i-1} \setminus \{W_i\}$. In addition, ∂U_1 is disjoint from A_i , other than at p_2 , and is disjoint from $\partial W'_i$. These conditions will shortly allow us to apply Claim 8.2 to U_1 .

Let F' denote the result of performing the finger move above to F . Note that

$$(8-1) \quad t_{\text{alt}}(F, \mathcal{U}_{i-1}) = t_{\text{alt}}(F', (\mathcal{U}_{i-1} \setminus \{W_i\}) \cup \{W'_i, U_1\})$$

by construction. Let V'_i denote the Whitney disc obtained as the union of V_i , W'_i and U_2 , as in Figure 17. Observe that the Whitney discs U_1 and V'_i pair the same double points, namely q_2 and p_2 . Consider the two collections of Whitney discs $(\mathcal{U}_{i-1} \setminus \{W_i\}) \cup \{W'_i, U_1\}$ and $(\mathcal{U}_{i-1} \setminus \{W_i\}) \cup \{W'_i, V'_i\}$ for the double points of F' . The two collections differ only in that one contains the disc U_1 and the other the disc V'_i . We will apply Claim 8.2 to change between the two collections. This is permitted since U_1 is framed, embedded, has interior disjoint from F' and the Whitney discs $(\mathcal{U}_{i-1} \setminus \{W_i\}) \cup \{W'_i\}$, and ∂U_1 is disjoint, other than at the endpoints, from the Whitney arcs of $(\mathcal{U}_{i-1} \setminus \{W_i\}) \cup \{W'_i, V'_i\}$, given by $A_i \cup \partial W'_i \cup \partial U_2$.

So, by Claim 8.2,

$$(8-2) \quad t_{\text{alt}}(F', (\mathcal{U}_{i-1} \setminus \{W_i\}) \cup \{W'_i, U_1\}) = t_{\text{alt}}(F', (\mathcal{U}_{i-1} \setminus \{W_i\}) \cup \{W'_i, V'_i\}).$$

By the proof of Claim 8.1 (see Figure 17),

$$(8-3) \quad t_{\text{alt}}(F', (\mathcal{U}_{i-1} \setminus \{W_i\}) \cup \{W'_i, V'_i\}) = t_{\text{alt}}(F', (\mathcal{U}_{i-1} \setminus \{W_i\}) \cup \{U_2, V_i\}).$$

Since U_2 is trivial, we can use it to undo the Whitney move, and obtain

$$(8-4) \quad t_{\text{alt}}(F', (\mathcal{U}_{i-1} \setminus \{W_i\}) \cup \{U_2, V_i\}) = t_{\text{alt}}(F, (\mathcal{U}_{i-1} \setminus \{W_i\}) \cup \{V_i\}) = t_{\text{alt}}(F, \mathcal{U}_i).$$

The combination of (8-1), (8-2), (8-3) and (8-4) implies $t_{\text{alt}}(F, \mathcal{U}_{i-1}) = t_{\text{alt}}(F, \mathcal{U}_i)$. This completes the proof that t_{alt} is independent of the choices of Whitney discs, and therefore completes the proof that t_{alt} is well defined.

Finally, by Lemma 7.6, we know that $t_{\text{alt}}(F, \mathcal{W}) = t(F, \mathcal{W})$ for \mathcal{W} a convenient collection, so t is well defined for convenient collections \mathcal{W} , as desired. □

Next we recall the statement of Theorem 1.6 for the convenience of the reader.

Theorem 1.6 *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1. Suppose that $\mu(F) = 0$ and that F has algebraically dual spheres. If F^∞ is not b -characteristic, then $\text{km}(F) = 0$. If F^∞ is b -characteristic, then the secondary embedding obstruction satisfies*

$$\text{km}(F) = t(F^\infty, \mathcal{W}^\infty) \in \mathbb{Z}/2$$

for every convenient collection of Whitney discs \mathcal{W}^∞ pairing all the double points of F^∞ .

Proof First we show that, if F^∞ is not b -characteristic then $\text{km}(F) = 0$. By Lemma 5.11, we reduce to the case that $\lambda_{\Sigma^\infty}|_{\partial\mathcal{B}(F^\infty)}$ is trivial, which implies that the function Θ is well defined on $\mathcal{B}(F^\infty)$. Since we assumed that F^∞ is not b -characteristic, there exists $B \in \mathcal{B}(F^\infty)$ such that $\Theta(B) = 1$, so we can apply Construction 7.2 and Lemma 7.3 to find a collection of Whitney discs \mathcal{W}^∞ for the double points of F^∞ with $t(F^\infty, \mathcal{W}^\infty) = 0$. Then, by Lemma 5.4, we know that $\text{km}(F) = 0$.

By Theorem 1.9, if F^∞ is b -characteristic, then $t(F^\infty, \mathcal{W}^\infty)$ is well defined, ie is independent of \mathcal{W}^∞ . As in Theorem 1.9, we denote the resulting invariant $t(F^\infty)$. We need to show that $\text{km}(F) = t(F^\infty)$.

Recall that b -characteristic implies r -characteristic by Lemma 5.18, and also r -characteristic implies s -characteristic by Remark 5.6. Therefore Lemma 5.4 applies, which says that, if $t(F^\infty) = t(F^\infty, \mathcal{W}^\infty) = 0$, then $\text{km}(F) = 0$. On the other hand, if $\text{km}(F) = 0$, then after a regular homotopy the double points of F can be paired up by a convenient collection of Whitney discs with interiors disjoint from F . Using these Whitney discs to calculate $t(F^\infty)$, and regular homotopy invariance of t from Theorem 1.9, it follows that $t(F^\infty) = 0$.

Thus we have shown that, for F^∞ b -characteristic and F with algebraically dual spheres, $\text{km}(F) = 0$ if and only if $t(F^\infty) = 0$ or, equivalently, $\text{km}(F) = t(F^\infty) \in \mathbb{Z}/2$, as desired. \square

9 Examples and applications

Proposition 9.1 *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1 and assume that $\mu(F) = 0$. If there are two orientation-preserving immersed loops in Σ that intersect transversely in an odd number of points and are null-homotopic in M , then F is not b -characteristic.*

Proof The two immersed loops in Σ from the assumption bound immersed discs in M . These discs give classes in $\mathcal{B}(F)$ by Construction 5.10 and, by assumption, $\lambda_{\Sigma}|_{\mathcal{B}(F)}$ is nontrivial. It follows by definition that F is not b -characteristic. \square

This applies to every simply connected target M whenever Σ has a component of positive genus. Proposition 9.1 also implies Corollaries 1.7 and 1.8, whose statements we recall, as follows.

Corollary 1.7 *If M is a simply connected 4-manifold and Σ is a connected, oriented surface with positive genus, then any generic immersion $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ with vanishing self-intersection number is not b -characteristic. Thus, if F has an algebraically dual sphere, then $\text{km}(F) = 0$, and, since $\pi_1(M)$ is good, the map F is regularly homotopic, relative to $\partial\Sigma$, to an embedding.*

Proof As $\pi_1(M)$ is trivial and Σ has positive genus, F is not b -characteristic by Proposition 9.1. By Theorem 1.6, it follows that $\text{km}(F) = 0$ if F has an algebraically dual sphere. In this case F is regularly homotopic, relative to $\partial\Sigma$, to an embedding, by Theorem 1.2. The theorem applies because $\pi_1(M)$ is good. \square

Corollary 1.8 *Let $F: (\Sigma, \partial\Sigma) \looparrowright (M, \partial M)$ be as in Convention 1.1, with $\mu(F) = 0$ and Σ connected. If F' is obtained from F by an ambient connected sum with an embedding $S^1 \times S^1 \hookrightarrow S^4$, then F' is not b -characteristic. Thus, if F has an algebraically dual sphere, then $\text{km}(F') = 0$, and, if $\pi_1(M)$ is good, then F' is regularly homotopic, relative to $\partial\Sigma$, to an embedding.*

Proof Since F' is obtained from F by an ambient connected sum with an embedding $S^1 \times S^1 \hookrightarrow S^4$, we can apply Proposition 9.1 to see that F' is not b -characteristic. By Theorem 1.6, it follows that $\text{km}(F') = 0$ if F has an algebraically dual sphere, as this sphere remains algebraically dual to F' . If in addition $\pi_1(M)$ is good, then, by Theorem 1.2, F is regularly homotopic, relative to $\partial\Sigma$, to an embedding. \square

Example 9.2 To illustrate the difference between r -characteristic and b -characteristic surfaces, we give an example of a surface that is r -characteristic but not b -characteristic. Consider any r -characteristic immersed sphere with trivial self-intersection number. Add a single trivial tube to obtain an immersed torus. As this will not change the intersection number with any closed surface, the new torus is still r -characteristic. But it fails to be b -characteristic by Corollary 1.8.

Example 9.3 We explain next why our methods allow us to obtain embeddings where [Freedman and Quinn 1990, Theorem 10.5A(1)] would not produce them (see the discussion directly following Theorem 1.2).

Let $f: S^2 \looparrowright M$ be a generic immersion in a 4-manifold with $\pi_1(M)$ good, equipped with an algebraically dual sphere and with $\text{km}(f) = 1$, for example a sphere representing a generator of $H_2(*\mathbb{C}\mathbb{P}^2)$. Other such spheres may be constructed as in [Kasprowski et al. 2021a, Theorem 2]. Let T be a generic immersion of a torus produced by adding a trivial tube to f , ie by taking the ambient connected sum of f with the standard embedding $S^1 \times S^1 \hookrightarrow S^4$. Then by Corollary 1.8 we see that $\text{km}(T) = 0$. Thus T is regularly homotopic to an embedding since $\pi_1(M)$ is good. Fix a 1-skeleton Σ_0 for $S^1 \times S^1$. Then T is not regularly homotopic to an embedding relative to Σ_0 , since the Kervaire–Milnor invariant for f restricted to the 2-cell(s) $(S^1 \times S^1) \setminus \nu\Sigma_0$ considered as a map to $M \setminus T(\nu\Sigma_0)$ equals $\text{km}(f) = 1$.

We emphasise that this holds for every choice of 1-skeleton $\Sigma_0 \subseteq S^1 \times S^1$. In order to apply the strategy of [Freedman and Quinn 1990, Theorem 10.5A(1)] to find an embedding, one needs to first make a

judicious choice of finger moves. But, without our theory, there is no clear strategy for finding these finger moves. To obtain an embedding obstruction in this way, matters are worse, since one would need to compute the Kervaire–Milnor invariant of the 2–skeleton for every choice of finger moves and for every choice of 1–skeleton.

Example 9.4 We construct an immersed torus with nontrivial Kervaire–Milnor invariant. In contrast to Proposition 9.1, the torus in this example is not π_1 –trivial. Consider an immersion f_1 of a 2–sphere in $*\mathbb{C}\mathbb{P}^2$ representing a generator of $H_2(*\mathbb{C}\mathbb{P}^2; \mathbb{Z})$ with trivial self-intersection number. Let $K: S^1 \hookrightarrow S^3$ be an arbitrary knot and consider the embedding of a torus given by the product $f_2 := K \times \text{Id}: S^1 \times S^1 \hookrightarrow S^3 \times S^1$. Let F denote the interior connected sum $f_1 \# f_2: S^1 \times S^1 \hookrightarrow W := *\mathbb{C}\mathbb{P}^2 \# (S^3 \times S^1)$.

First we claim that F is b –characteristic. To see this, we start by computing $H_2(W, S^1 \times S^1; \mathbb{Z}/2)$ using the long exact sequence of the pair with $\mathbb{Z}/2$ coefficients:

$$\begin{array}{ccccccc}
 H_2(S^1 \times S^1) & \twoheadrightarrow & H_2(W) & \xrightarrow{0} & H_2(W, S^1 \times S^1) & \longrightarrow & H_1(S^1 \times S^1) & \twoheadrightarrow & H_1(W) \\
 & & & & & & \cong & & \cong \\
 & & & & & & \mathbb{Z}/2 \oplus \mathbb{Z}/2 & & \mathbb{Z}/2
 \end{array}$$

Therefore, $H_2(W, S^1 \times S^1; \mathbb{Z}/2) \cong \mathbb{Z}/2$ is generated by $S \times \{p\}$, where $S \subseteq S^3$ is a Seifert surface for the knot $K(S^1)$ and $p \in S^1$. The intersection form of $S^1 \times S^1$ restricted to ∂S is trivial. Since Θ is well defined on homology classes we can compute it using S . But S has interior disjoint from the image of F , embedded boundary and trivial relative Euler number, so $\Theta(S) = 0$. It follows that Θ vanishes on all of $H_2(W, S^1 \times S^1; \mathbb{Z}/2)$, in particular it vanishes on the subset $\mathcal{B}(F)$. Thus F is b –characteristic, as claimed.

Observe that $\text{km}(f_1) = 1$ inside $*\mathbb{C}\mathbb{P}^2$ (see eg [Freedman and Quinn 1990, Section 10.8]). We can pick a convenient collection of Whitney discs for f_1 in $*\mathbb{C}\mathbb{P}^2$. Since f_2 is an embedding, these constitute a convenient collection of Whitney discs for F . It follows that $\text{km}(F) = \text{km}(f_1) = 1$. Note that the choice of knot K was irrelevant, since, for any two choices, the resulting immersions F are regularly homotopic and hence have equal Kervaire–Milnor invariant.

Example 9.5 In the previous example we constructed a generically immersed torus in $*\mathbb{C}\mathbb{P}^2 \# (S^1 \times S^3)$ with nontrivial Kervaire–Milnor invariant. In particular, this torus is not homotopic to an embedding (see Section 1.4). Now we show that, in contrast to this, every map f from a closed surface Σ to $S^1 \times S^3$ is homotopic to an embedding. Note that these classes do not have algebraically dual spheres since $\pi_2(S^1 \times S^3) = 0$. The surfaces in the regular homotopy class with $\mu(f)_1 = 0$ are either not b –characteristic or $t(f^\infty)$ vanishes.

Since the projection $S^1 \times S^3 \rightarrow S^1$ is 3–connected, the induced map $[\Sigma, S^1 \times S^3] \rightarrow [\Sigma, S^1]$ is bijective. In particular, the homotopy class of a map $f: \Sigma \rightarrow S^1 \times S^3$ is determined by the induced map on fundamental groups.

We first consider the case that Σ is connected. Since $\pi_1(S^1 \times S^3) \cong \mathbb{Z}$, we can find a generating set for $\pi_1(\Sigma)$ such that at most one generator is nontrivial in $\pi_1(S^1 \times S^3)$. Thus there exists a decomposition $\Sigma = H \# \Sigma'$, where H is either a sphere, a torus or a Klein bottle, with respect to which f can be written as an internal connected sum $T \# f'$, where T is a map on H and f' is π_1 -trivial. In particular, f' is homotopic to an embedding inside a ball $D^4 \subseteq S^1 \times S^3$. It remains to show that T is homotopic to an embedding, which will show that the connected sum is homotopic to an embedding.

If H is a sphere, we are done. If H is a torus, let $i : S^1 \times S^1 \hookrightarrow S^3$ be an embedding. For each $k \in \mathbb{Z}$, define the embedding $h'_k : S^1 \times S^1 \rightarrow S^1 \times (S^1 \times S^1)$ by $(s, t) \mapsto (s^k, (s, t))$. Let $h_k := (\text{Id} \times i) \circ h'_k$. There exists some k and some identification of H with $S^1 \times S^1$ such that T and h_k induce the same map on fundamental groups and thus are homotopic. If H is a Klein bottle, let $p : H \rightarrow S^1$ be a fibre bundle with fibre S^1 . For each $k \in \mathbb{Z}$, there exists an immersion $i : H \looparrowright S^3$ such that $h_k(x) := (p(x)^k, i(x))$ is an embedding $H \hookrightarrow S^1 \times S^3$. As before, there exists some k such that T and h_k are homotopic.

The above embeddings can be realised as embeddings into $S^1 \times D^3 \subseteq S^1 \times S^3$. The argument generalises to disconnected surfaces by picking disjoint copies of $S^1 \times D^3$ in $S^1 \times S^3$ for each connected component.

Next we prove Proposition 1.10 and Corollary 1.11, which we restate for the convenience of the reader.

Proposition 1.10 *Let M_1 and M_2 be oriented 4-manifolds. Let $F_1 : (\Sigma_1, \partial\Sigma_1) \looparrowright (M_1, \partial M_1)$ and $F_2 : (\Sigma_2, \partial\Sigma_2) \looparrowright (M_2, \partial M_2)$ be generic immersions of connected, compact, oriented surfaces, each with vanishing self-intersection number. If F_i is b -characteristic for each i , then both the disjoint union*

$$F_1 \sqcup F_2 : \Sigma_1 \sqcup \Sigma_2 \looparrowright M_1 \# M_2$$

and any interior connected sum

$$F_1 \# F_2 : \Sigma_1 \# \Sigma_2 \looparrowright M_1 \# M_2$$

are b -characteristic, and satisfy

$$t(F_1 \sqcup F_2) = t(F_1 \# F_2) = t(F_1) + t(F_2).$$

Proof The vanishing of the self-intersection number of F_i is witnessed by a convenient collection of Whitney discs \mathcal{W}_i in M_i for each i . The union $\mathcal{W}_1 \sqcup \mathcal{W}_2$, now considered in $M_1 \# M_2$, shows that the intersection and self-intersection numbers of $F_1 \sqcup F_2$, as well as for $F_1 \# F_2$, vanish in $M_1 \# M_2$. Since the union $\mathcal{W}_1 \sqcup \mathcal{W}_2$ pairs all the double points of $F_1 \sqcup F_2$ (resp. $F_1 \# F_2$), and since components of \mathcal{W}_i cannot intersect $F_j(\Sigma_j)$ for all $i \neq j$, the claimed relationship $t(F_1 \sqcup F_2) = t(F_1 \# F_2) = t(F_1) + t(F_2)$ holds as long as $F_1 \sqcup F_2$ and $F \# F_2$ are b -characteristic.

As a preliminary step, note that neither F_i has a framed dual sphere in M_i , since otherwise it would not be s -characteristic, and therefore not b -characteristic by Lemma 5.18. As a result, $\Sigma_i^{\circ} = \Sigma_i$ for $i = 1, 2$.

Next we consider the immersion $F_1 \sqcup F_2 : (\Sigma_1 \sqcup \Sigma_2, \partial\Sigma_1 \sqcup \partial\Sigma_2) \looparrowright M_1 \# M_2$. Let $S \subseteq M_1 \# M_2$ denote a connected sum 3-sphere. Consider a band $[B] \in H_2(M_1 \# M_2, \Sigma_1 \sqcup \Sigma_2)$. By (topological) transversality, we can assume that B is immersed, the double points of B are disjoint from S , and the

intersection $B \cap S$ corresponds to an embedded 1–manifold in the interior of the domain of B , since $\partial B \subseteq \Sigma_1 \sqcup \Sigma_2 \subseteq (M_1 \# M_2) \setminus S$. The image of this 1–manifold in S is embedded and bounds a collection of immersed (perhaps intersecting) discs in S . Surger B using two copies each of these discs to produce $B_1 \subseteq M_1$ and $B_2 \subseteq M_2$, where each B_i is an immersed collection of surfaces with $\partial B_i \subseteq \Sigma_i$.

Each component of B_i can be replaced by a band as follows. Recall that since each M_i and Σ_i is oriented, there is no condition on Stiefel–Whitney classes for bands, and we need only arrange that each component is either a Möbius band or an annulus. By considering the Euler characteristic, we see that each component is homeomorphic to either a sphere, an $\mathbb{R}P^2$, a disc, a Möbius band or an annulus. Then use the tubing procedure from Construction 5.10 to replace each sphere, $\mathbb{R}P^2$ or disc component by a band. More precisely, choose a small disc on Σ_1 or Σ_2 , as appropriate, away from all Whitney arcs and double points, and tube into the disc, sphere or $\mathbb{R}P^2$.

Since each F_i is b –characteristic, $\lambda_{\Sigma_i}|_{\partial \mathcal{B}(F_i)}$ is trivial for each i . Therefore, $\lambda_{\Sigma_1 \sqcup \Sigma_2}$ is trivial on $\partial B = \partial B_1 \cup \partial B_2$. It follows by Lemma 5.16(iv) that $\Theta: \mathcal{B}(F_1 \sqcup F_2) \rightarrow \mathbb{Z}/2$ is well defined. By Lemma 7.8, Θ extends to a linear map $\langle \mathcal{B}(F_1 \sqcup F_2) \rangle \rightarrow \mathbb{Z}/2$ on the subspace

$$\langle \mathcal{B}(F_1 \sqcup F_2) \rangle \subseteq H_2(M_1 \# M_2, \Sigma_1 \sqcup \Sigma_2; \mathbb{Z}/2)$$

generated by the bands. Then, since $[B_1 \cup B_2] = [B] \in H_2(M_1 \# M_2, \Sigma_1 \sqcup \Sigma_2; \mathbb{Z}/2)$, we see that $\Theta(B) = \Theta(B_1) + \Theta(B_2)$.

For each i , the value of $\Theta(B_i)$ does not depend on whether the ambient manifold is M_i or $M_1 \# M_2$, since B_i does not intersect $F_j(\Sigma_j)$ for all $i \neq j$ (see Definition 5.12). Since each F_i is b –characteristic, $\Theta(B) = \Theta(B_1) + \Theta(B_2) = 0 + 0 = 0 \in \mathbb{Z}/2$. This completes the proof that $F_1 \sqcup F_2$ is b –characteristic.

Now we consider the connected sum $F_1 \# F_2$. Let $B \in H_2(M_1 \# M_2, \Sigma_1 \# \Sigma_2)$ be a band. As above, we assume that the intersection $B \cap S$ corresponds to an embedded 1–manifold in the domain of B . Unlike above, this may include embedded arcs with endpoints on the boundary. These endpoints are mapped to the intersection $(F_1 \# F_2)(\Sigma_1 \# \Sigma_2) \cap S$. By connecting the endpoints with arcs on $(F_1 \# F_2)(\Sigma_1 \# \Sigma_2) \cap S$, we again get a collection of closed circles in S , which bound an immersed collection of discs in S . Surger using these discs as before to produce collections $B_i \subseteq M_i$. Once again, each component of B_i is homeomorphic to either a sphere, an $\mathbb{R}P^2$, a disc, a Möbius band or an annulus. By Construction 5.10 applied to the sphere, $\mathbb{R}P^2$ and disc components, we may arrange that each component is a band. The argument of the previous paragraph now applies to show that $F_1 \# F_2$ is b –characteristic. \square

Corollary 1.11 *For any g , there exists a smooth, closed 4–manifold M_g , a closed, connected, oriented surface Σ_g of genus g , and a smooth, b –characteristic, generic immersion $F: \Sigma_g \looparrowright M_g$ with $t(F) \neq 0$ and therefore $\text{km}(F) \neq 0$.*

Proof By the same proof as in Example 9.4, for any knot K the product $T := K \times \text{Id}: S^1 \times S^1 \rightarrow S^3 \times S^1$ is an embedded b –characteristic torus. Since T is an embedding, $t(T) = 0$. A computation using the intersection form shows that a generic immersion $S: S^2 \rightarrow \mathbb{C}P^2$ representing three times a generator of

$H_2(\mathbb{C}\mathbb{P}^2; \mathbb{Z})$ is s -characteristic. Since $\pi_1(\mathbb{C}\mathbb{P}^2)$ has no 2-torsion, the map S is also r -characteristic and thus b -characteristic by Lemma 5.18. We will show that $\text{km}(S) = 1$. This was the original example of Kervaire and Milnor [1961]. To see that $\text{km}(S) = 1$, represent S in the following way. Take a cuspidal cubic, which is a smooth embedding of a 2-sphere away from a single singular point. In a neighbourhood of the singular point we see a cone on the trefoil. Replace a neighbourhood of the singular point with an immersed disc Δ in D^4 with boundary the trefoil, and two double points that are paired by a framed Whitney disc that intersects Δ once. Alternatively, we can compute $t(S)$ as $\frac{1}{8}(\sigma(\mathbb{C}\mathbb{P}^2) - S \cdot S) = \frac{1}{8}(1 - 9) \equiv 1 \pmod{2}$; see Section 3. This gives us the case $g = 0$. Next, by Proposition 1.10, for every $g \in \mathbb{N}$,

$$S \# \#^g T : \Sigma_g \rightarrow \mathbb{C}\mathbb{P}^2 \# \#^g (S^3 \times S^1)$$

is a b -characteristic generic immersion of a closed surface of genus g with nontrivial t , and therefore $\text{km}(S \# \#^g T) \neq 0$. In particular, $S \# \#^g T$ is not regularly homotopic to an embedding. Note these examples are smooth, but have no algebraically dual sphere. We could replace $(\mathbb{C}\mathbb{P}^2, S)$ with $(\mathbb{C}\mathbb{P}^2 \# \#^8 \mathbb{C}\mathbb{P}^2, S')$, where S' is a generic immersion representing the class $(3, 1, \dots, 1) \in \mathbb{Z}^9 \cong H_2(\mathbb{C}\mathbb{P}^2 \# \#^8 \mathbb{C}\mathbb{P}^2)$, to obtain an example with an algebraically dual sphere and $\text{km}(S') = \frac{1}{8}(-7 - 1) \equiv 1 \pmod{2}$. □

Remark 9.6 Let M denote the infinite connected sum $\mathbb{C}\mathbb{P}^2 \# \#^\infty (S^3 \times S^1)$. The proof of Corollary 1.11, along with the formula from Proposition 1.10, shows that, for every g , there exists a smooth generic immersion $F : \Sigma_g \looparrowright M$ with $t(F) \neq 0$ and therefore $\text{km}(F) \neq 0$. The following proposition shows that, if there is such a compact 4-manifold M and such an F , then the 4-manifolds must have nonabelian fundamental group. In other words, if there is an immersed surface in a 4-manifold with abelian fundamental group with nontrivial km , then we give a bound on the complexity of that surface.

Proposition 9.7 *Let M be a compact 4-manifold such that $\pi_1(M)$ is abelian with n generators. Let $F : \Sigma \looparrowright M$ be a b -characteristic generic immersion, where Σ is a closed, connected surface. Then the Euler characteristic satisfies $\chi(\Sigma) \geq -2n$.*

Proof Suppose that $\chi(\Sigma) < -2n$. Note that Σ can be written as a connected sum of a genus g orientable surface Σ' for some $g > n$ with zero, one or two copies of $\mathbb{R}\mathbb{P}^2$. There exists a surjection $\mathbb{Z}^n \twoheadrightarrow \pi_1(M)$. Then the induced map $H_1(\Sigma') \rightarrow H_1(M) \cong \pi_1(M)$ admits a lift $H_1(\Sigma') \rightarrow \mathbb{Z}^n$, which has kernel of rank at least $2g - n > g$. So there exist closed curves γ_1 and γ_2 in $\Sigma' \setminus D^2 \subseteq \Sigma$ that are null-homotopic in M and $\lambda_\Sigma(\gamma_1, \gamma_2) \equiv 1 \pmod{2}$. It follows that F is not b -characteristic. □

Next we prove our corollaries on knot theory from Section 1.5.

Corollary 1.15 *For every knot $K \subseteq S^3$,*

- (1) $g_M(K) = 0$ for every simply connected 4-manifold M not homeomorphic to one of $S^4, \mathbb{C}\mathbb{P}^2$ or $*\mathbb{C}\mathbb{P}^2$;

- (2) $g_{\mathbb{C}\mathbb{P}^2}(K) \leq 1$ and $g_{\mathbb{C}\mathbb{P}^2}(\#^3 T(2, 3)) = 1$; and
 (3) $g_{*\mathbb{C}\mathbb{P}^2}(K) \leq 1$ and $g_{*\mathbb{C}\mathbb{P}^2}(\#^2 T(2, 3)) = 1$.

Proof Let $K \subseteq S^3$ be an arbitrary knot and let M be an arbitrary closed, simply connected 4–manifold. Let Δ' be a generically immersed disc bounded by K in a collar $S^3 \times [0, 1]$ of ∂M° . Since M is simply connected, every class in $H_2(M; \mathbb{Z}) \cong \pi_2(M)$ is represented by a generically immersed sphere. By assumption, M is not homeomorphic to S^4 and thus $H_2(M; \mathbb{Z})$ is nontrivial. Since M is closed, every primitive class $\alpha \in H_2(M; \mathbb{Z})$ has an algebraic dual $\beta \in H_2(M; \mathbb{Z})$, ie $\lambda(\alpha, \beta) = 1$. Represent α and β by generically immersed spheres, and tube the interior of Δ' into β to obtain Δ . Add local cusps to arrange $\mu(\Delta) = 0$.

First we prove (1). In this case we claim that in the construction of Δ we can choose the primitive class α to satisfy $\lambda(\alpha, \alpha) \in 2\mathbb{Z}$, as we explain presently. Then the disc Δ constructed above is not r –characteristic, since $\Delta \cdot \alpha \not\equiv \alpha \cdot \alpha \pmod{2}$ (see Remark 5.6). By Theorem 5.7, this implies that $\text{km}(\Delta) = 0$. Since the disc Δ has the algebraically dual sphere α and $\pi_1(M) = 1$ is good, by Theorem 1.2, Δ is homotopic rel boundary to an embedding. To see the claim regarding α , note that, when $M \not\cong S^4, \mathbb{C}\mathbb{P}^2, *\mathbb{C}\mathbb{P}^2$, the group $H_2(M; \mathbb{Z})$ has rank at least 2 by the classification of closed, simply connected 4–manifolds up to homeomorphism. Then $H_2(M; \mathbb{Z})$ has a summand isomorphic to $\mathbb{Z} \oplus \mathbb{Z}$, so the classes x, y and $x + y$, for the generators x and y of the \mathbb{Z} factors, are primitive, and at least one of $\lambda(x, x), \lambda(y, y)$ or $\lambda(x + y, x + y)$ is even.

In (2) and (3), we have $M = \mathbb{C}\mathbb{P}^2$ or $*\mathbb{C}\mathbb{P}^2$. The only primitive classes are $\pm[\mathbb{C}\mathbb{P}^1]$, so we choose $\alpha = \beta = [\mathbb{C}\mathbb{P}^1]$ in the construction of the first paragraph. We construct the disc Δ as before, but are no longer able to conclude that it is r –characteristic. However, by Corollary 1.8, we know that the connected sum of Δ with an unknotted torus is homotopic to an embedding. This completes the proof of the first parts of (2) and (3).

Now we prove that $g_{\mathbb{C}\mathbb{P}^2}(\#^3 T(2, 3)) = 1$. Let $K := \#^3 T(2, 3)$. Let $g_{\mathbb{C}\mathbb{P}^2}^d(K)$ denote the minimal genus of a surface bounded by K in $(\mathbb{C}\mathbb{P}^2)^\circ$ in the homology class $d \in \mathbb{Z} \cong H_2(\mathbb{C}\mathbb{P}^2; \mathbb{Z})$. First we consider $d = \pm 1$, where the class is b –characteristic (or equivalently, s –characteristic; see Lemma 5.18). As before, construct the disc $\Delta' \subseteq S^3 \times [0, 1]$, and tube into $\mathbb{C}\mathbb{P}^1$ to obtain the disc Δ . We assume that Δ' has trivial self-intersection number, so $1 = \text{Arf}(K) = t(\Delta')$ by [Matsumoto 1978; Freedman and Kirby 1978; Conant et al. 2014, Lemma 10]. Since $\mathbb{C}\mathbb{P}^1$ is embedded disjointly from Δ' , $t(\Delta) = 1$. Thus, by Theorem 1.9, Δ is not homotopic to an embedding and so $g_{\mathbb{C}\mathbb{P}^2}^{\pm 1}(K) \neq 0$.

Next let $\sigma_d(K) := \sigma_K(e^{\pi i(d-1)/d})$, where σ_K denotes the Levine–Tristram signature function of K . By [Gilmer 1981; Viro 1975], for even d ,

$$2g_{\mathbb{C}\mathbb{P}^2}^d(K) + 1 \geq \left| \frac{1}{2}d^2 - 1 - \sigma(K) \right|,$$

while, if d is divisible by an odd prime p , then

$$2g_{\mathbb{C}\mathbb{P}^2}^d(K) + 1 \geq \left| \frac{p^2 - 1}{2p^2} d^2 - 1 - \sigma_d(K) \right|.$$

In our case, $\sigma(K) = \sigma_d(K) = -6$ for all d , and so $g_{\mathbb{C}\mathbb{P}^2}^d(K) \geq 1$ for all $d \neq \pm 1$. This completes the argument that $g_{\mathbb{C}\mathbb{P}^2}(K) = 1$.

Finally we show that $g_{*\mathbb{C}\mathbb{P}^2}(\#^2 T(2, 3)) = 1$. Write $K := \#^2 T(2, 3)$ and let $g_{*\mathbb{C}\mathbb{P}^2}^d(K)$ denote the minimal genus of a surface bounded by K in $(*\mathbb{C}\mathbb{P}^2)^\circ$ in the homology class $d \in H_2(*\mathbb{C}\mathbb{P}^2; \mathbb{Z})$. For $d = \pm 1$, modify the argument above for the case of $\mathbb{C}\mathbb{P}^2$, using that tubing into a sphere representing a generator of $H_2(*\mathbb{C}\mathbb{P}^2; \mathbb{Z})$ to obtain a disc Δ' adds 1 to the t count, and so $1 = 1 + \text{Arf}(K) = t(\Delta')$. Therefore, again by Theorem 1.9, $g_{*\mathbb{C}\mathbb{P}^2}^{\pm 1}(K) \neq 0$. Next, for $*\mathbb{C}\mathbb{P}^2$, the same inequalities from [Gilmer 1981; Viro 1975] hold, and $\sigma(K) = \sigma_d(K) = -4$ for all d . Therefore, applying the inequalities, we see that $g_{*\mathbb{C}\mathbb{P}^2}^d(K) \geq 1$ for all d . It follows that $g_{*\mathbb{C}\mathbb{P}^2}(K) = 1$, as asserted. \square

Corollary 1.16 For any knot $K \subseteq S^3$, $g_{\pm 1}^{\text{sh}}(K) = \text{Arf}(K) \in \{0, 1\}$.

Proof A generator of $H_2(X_{\pm 1}(K); \mathbb{Z})$ can be represented by a generically immersed sphere F which is b -characteristic (or equivalently s -characteristic; see Lemma 5.18), has trivial $\mu(F)$, and has an algebraically dual sphere. We also recall from [Matsumoto 1978; Freedman and Kirby 1978; Conant et al. 2014, Lemma 10] that $\text{Arf}(K)$ coincides with the count $t(F)$. Then, by Theorems 1.6 and 1.9, the sphere F is homotopic to an embedding if and only if $\text{Arf}(K) = 0$. We have an embedded torus representative for both generators by Corollary 1.8. \square

References

- [Aït Nouh 2009] **M Aït Nouh**, *Genera and degrees of torus knots in $\mathbb{C}\mathbb{P}^2$* , J. Knot Theory Ramifications 18 (2009) 1299–1312 MR Zbl
- [Behrens et al. 2021] **S Behrens, B Kalmár, M H Kim, M Powell, A Ray** (editors), *The disc embedding theorem*, Oxford Univ. Press (2021) MR Zbl
- [Casson 1986] **A J Casson**, *Three lectures on new-infinite constructions in 4-dimensional manifolds*, from “À la recherche de la topologie perdue” (L Guillou, A Marin, editors), Progr. Math. 62, Birkhäuser, Boston, MA (1986) 201–244 MR Zbl
- [Cochran and Ray 2016] **T D Cochran, A Ray**, *Shake slice and shake concordant knots*, J. Topol. 9 (2016) 861–888 MR Zbl
- [Cochran et al. 2003] **T D Cochran, K E Orr, P Teichner**, *Knot concordance, Whitney towers and L^2 -signatures*, Ann. of Math. 157 (2003) 433–519 MR Zbl
- [Conant et al. 2012a] **J Conant, R Schneiderman, P Teichner**, *Universal quadratic forms and Whitney tower intersection invariants*, from “Proceedings of the Freedman Fest” (R Kirby, V Krushkal, Z Wang, editors), Geom. Topol. Monogr. 18, Geom. Topol. Publ., Coventry (2012) 35–60 MR Zbl

- [Conant et al. 2012b] **J Conant, R Schneiderman, P Teichner**, *Whitney tower concordance of classical links*, *Geom. Topol.* 16 (2012) 1419–1479 MR Zbl
- [Conant et al. 2014] **J Conant, R Schneiderman, P Teichner**, *Milnor invariants and twisted Whitney towers*, *J. Topol.* 7 (2014) 187–224 MR Zbl
- [Connelly 1971] **R Connelly**, *A new proof of Brown’s collaring theorem*, *Proc. Amer. Math. Soc.* 27 (1971) 180–182 MR Zbl
- [Feller et al. 2021] **P Feller, A N Miller, M Nagel, P Orson, M Powell, A Ray**, *Embedding spheres in knot traces*, *Compos. Math.* 157 (2021) 2242–2279 MR Zbl
- [Freedman 1982] **M H Freedman**, *The topology of four-dimensional manifolds*, *J. Differential Geom.* 17 (1982) 357–453 MR Zbl
- [Freedman and Kirby 1978] **M Freedman, R Kirby**, *A geometric proof of Rochlin’s theorem*, from “Algebraic and geometric topology, II” (R J Milgram, editor), *Proc. Sympos. Pure Math.* 32, Amer. Math. Soc., Providence, RI (1978) 85–97 MR Zbl
- [Freedman and Quinn 1990] **M H Freedman, F Quinn**, *Topology of 4-manifolds*, Princeton Math. Ser. 39, Princeton Univ. Press (1990) MR Zbl
- [Freedman and Teichner 1995] **M H Freedman, P Teichner**, *4-manifold topology, I: Subexponential groups*, *Invent. Math.* 122 (1995) 509–529 MR Zbl
- [Gilmer 1981] **P M Gilmer**, *Configurations of surfaces in 4-manifolds*, *Trans. Amer. Math. Soc.* 264 (1981) 353–380 MR Zbl
- [Golubitsky and Guillemin 1973] **M Golubitsky, V Guillemin**, *Stable mappings and their singularities*, Graduate Texts in Math. 14, Springer (1973) MR Zbl
- [Guillou and Marin 1980] **L Guillou, A Marin**, *Une extension d’un théorème de Rohlin sur la signature*, from “Seminar on real algebraic geometry” (J J Risler, editor), *Publ. Math. Univ. Paris VII* 9, Univ. Paris VII (1980) 69–80 MR Zbl
- [Hambleton and Kreck 1993] **I Hambleton, M Kreck**, *Cancellation of hyperbolic forms and topological four-manifolds*, *J. Reine Angew. Math.* 443 (1993) 21–47 MR Zbl
- [Hirsch 1976] **M W Hirsch**, *Differential topology*, Graduate Texts in Math. 33, Springer (1976) MR Zbl
- [Kasprowski et al. 2017] **D Kasprowski, M Land, M Powell, P Teichner**, *Stable classification of 4-manifolds with 3-manifold fundamental groups*, *J. Topol.* 10 (2017) 827–881 MR Zbl
- [Kasprowski et al. 2020] **D Kasprowski, M Powell, P Teichner**, *Algebraic criteria for stable diffeomorphism of spin 4-manifolds*, preprint (2020) arXiv 2006.06127
- [Kasprowski et al. 2021a] **D Kasprowski, P Lambert-Cole, M Land, A G Lecuona**, *Topologically flat embedded 2-spheres in specific simply connected 4-manifolds*, from “2019–20 MATRIX annals” (D R Wood, J de Gier, C E Praeger, T Tao, editors), *MATRIX Book Ser.* 4, Springer (2021) 111–116 MR Zbl
- [Kasprowski et al. 2021b] **D Kasprowski, M Powell, P Teichner**, *The Kervaire–Milnor invariant in the stable classification of spin 4-manifolds*, preprint (2021) arXiv 2105.12153
- [Kervaire and Milnor 1961] **M A Kervaire, J W Milnor**, *On 2-spheres in 4-manifolds*, *Proc. Nat. Acad. Sci. U.S.A.* 47 (1961) 1651–1657 MR Zbl
- [Kim et al. 2021] **M H Kim, P Orson, J Park, A Ray**, *Good groups*, from “The disc embedding theorem” (S Behrens, B Kalmár, M H Kim, M Powell, A Ray, editors), Oxford Univ. Press (2021) 273–282 MR Zbl

- [Kirby and Taylor 2001] **R C Kirby, L R Taylor**, *A survey of 4-manifolds through the eyes of surgery*, from “Surveys on surgery theory, II” (S Cappell, A Ranicki, J Rosenberg, editors), Ann. of Math. Stud. 149, Princeton Univ. Press (2001) 387–421 MR Zbl
- [Krushkal and Quinn 2000] **V S Krushkal, F Quinn**, *Subexponential groups in 4-manifold topology*, Geom. Topol. 4 (2000) 407–430 MR Zbl
- [Lee and Wilczyński 1990] **R Lee, D M Wilczyński**, *Locally flat 2-spheres in simply connected 4-manifolds*, Comment. Math. Helv. 65 (1990) 388–412 MR Zbl
- [Lee and Wilczyński 1997] **R Lee, D M Wilczyński**, *Representing homology classes by locally flat surfaces of minimum genus*, Amer. J. Math. 119 (1997) 1119–1137 MR Zbl
- [Lisca and Matić 1998] **P Lisca, G Matić**, *Stein 4-manifolds with boundary and contact structures*, Topology Appl. 88 (1998) 55–66 MR Zbl
- [Manolescu et al. 2024] **C Manolescu, M Marengon, L Piccirillo**, *Relative genus bounds in indefinite four-manifolds*, Math. Ann. (online publication January 2024)
- [Marengon et al. 2024] **M Marengon, A N Miller, A Ray, A I Stipsicz**, *A note on surfaces in $\mathbb{C}P^2$ and $\mathbb{C}P^2 \# \mathbb{C}P^2$* , Proc. Amer. Math. Soc. Ser. B 11 (2024) 187–199 MR Zbl
- [Massey 1969] **W S Massey**, *Proof of a conjecture of Whitney*, Pacific J. Math. 31 (1969) 143–156 MR Zbl
- [Matsumoto 1978] **Y Matsumoto**, *Secondary intersectional properties of 4-manifolds and Whitney’s trick*, from “Algebraic and geometric topology, II” (R J Milgram, editor), Proc. Sympos. Pure Math. 32, Amer. Math. Soc., Providence, RI (1978) 99–107 MR Zbl
- [Matsumoto 1986] **Y Matsumoto**, *An elementary proof of Rochlin’s signature theorem and its extension by Guillou and Marin*, from “À la recherche de la topologie perdue” (L Guillou, A Marin, editors), Progr. Math. 62, Birkhäuser, Boston, MA (1986) 119–139 MR Zbl
- [Norman 1969] **R A Norman**, *Dehn’s lemma for certain 4-manifolds*, Invent. Math. 7 (1969) 143–147 MR Zbl
- [Pichelmeyer 2020] **J Pichelmeyer**, *Genera of knots in the complex projective plane*, J. Knot Theory Ramifications 29 (2020) art. id. 2050081 MR Zbl
- [Powell and Ray 2021a] **M Powell, A Ray**, *Basic geometric constructions*, from “The disc embedding theorem” (S Behrens, B Kalmár, M H Kim, M Powell, A Ray, editors), Oxford Univ. Press (2021) 217–226 MR
- [Powell and Ray 2021b] **M Powell, A Ray**, *Intersection numbers and the statement of the disc embedding theorem*, from “The disc embedding theorem” (S Behrens, B Kalmár, M H Kim, M Powell, A Ray, editors), Oxford Univ. Press (2021) 155–170 MR Zbl
- [Powell et al. 2020] **M Powell, A Ray, P Teichner**, *The 4-dimensional disc embedding theorem and dual spheres*, preprint (2020) arXiv 2006.05209
- [Quinn 1982] **F Quinn**, *Ends of maps, III: Dimensions 4 and 5*, J. Differential Geom. 17 (1982) 503–521 MR Zbl
- [Rokhlin 1952] **V A Rokhlin**, *New results in the theory of four-dimensional manifolds*, Doklady Akad. Nauk SSSR 84 (1952) 221–224 MR Zbl In Russian
- [Rokhlin 1972] **V A Rokhlin**, *Proof of Gudkov’s hypothesis*, Funct. Anal. Appl. 6 (1972) 136–138 Zbl
- [Schneiderman 2003] **R Schneiderman**, *Algebraic linking numbers of knots in 3-manifolds*, Algebr. Geom. Topol. 3 (2003) 921–968 MR Zbl
- [Schneiderman and Teichner 2001] **R Schneiderman, P Teichner**, *Higher order intersection numbers of 2-spheres in 4-manifolds*, Algebr. Geom. Topol. 1 (2001) 1–29 MR Zbl

- [Stong 1994] **R Stong**, *Existence of π_1 -negligible embeddings in 4-manifolds: a correction to Theorem 10.5 of [Freedman and Quinn 1990]*, Proc. Amer. Math. Soc. 120 (1994) 1309–1314 MR Zbl
- [Tristram 1969] **A G Tristram**, *Some cobordism invariants for links*, Proc. Cambridge Philos. Soc. 66 (1969) 251–264 MR Zbl
- [Viro 1975] **O Y Viro**, *Positioning in codimension 2, and the boundary*, Uspehi Mat. Nauk 30 (1975) 231–232 MR Zbl In Russian
- [Wall 1970] **C T C Wall**, *Surgery on compact manifolds*, Lond. Math. Soc. Monogr. 1, Academic, London (1970) MR Zbl
- [Yasuhara 1991] **A Yasuhara**, *(2, 15)-torus knot is not slice in $\mathbb{C}P^2$* , Proc. Japan Acad. Ser. A Math. Sci. 67 (1991) 353–355 MR Zbl
- [Yasuhara 1992] **A Yasuhara**, *On slice knots in the complex projective plane*, Rev. Mat. Univ. Complut. Madrid 5 (1992) 255–276 MR Zbl

*School of Mathematical Sciences, University of Southampton
Southampton, United Kingdom*

*School of Mathematics and Statistics, University of Glasgow
Glasgow, United Kingdom*

*Max Planck Institute for Mathematics
Bonn, Germany*

*Max Planck Institute for Mathematics
Bonn, Germany*

d.kasprowski@soton.ac.uk, mark.powell@glasgow.ac.uk, aruray@mpim-bonn.mpg.de,
teichner@mpim-bonn.mpg.de

Proposed: Ciprian Manolescu
Seconded: Anna Wienhard, Mladen Bestvina

Received: 4 October 2022
Revised: 28 February 2023

GEOMETRY & TOPOLOGY

msp.org/gt

MANAGING EDITOR

András I Stipsicz Alfréd Rényi Institute of Mathematics
stipsicz@renyi.hu

BOARD OF EDITORS

Mohammed Abouzaid	Stanford University abouzaid@stanford.edu	Mark Gross	University of Cambridge mgross@dpmms.cam.ac.uk
Dan Abramovich	Brown University dan_abramovich@brown.edu	Rob Kirby	University of California, Berkeley kirby@math.berkeley.edu
Ian Agol	University of California, Berkeley ianagol@math.berkeley.edu	Bruce Kleiner	NYU, Courant Institute bkleiner@cims.nyu.edu
Arend Bayer	University of Edinburgh arend.bayer@ed.ac.uk	Sándor Kovács	University of Washington skovacs@uw.edu
Mark Behrens	University of Notre Dame mbehren1@nd.edu	Urs Lang	ETH Zürich urs.lang@math.ethz.ch
Mladen Bestvina	University of Utah bestvina@math.utah.edu	Marc Levine	Universität Duisburg-Essen marc.levine@uni-due.de
Martin R Bridson	University of Oxford bridson@maths.ox.ac.uk	Ciprian Manolescu	University of California, Los Angeles cm@math.ucla.edu
Jim Bryan	University of British Columbia jbryan@math.ubc.ca	Haynes Miller	Massachusetts Institute of Technology hrm@math.mit.edu
Dmitri Burago	Pennsylvania State University burago@math.psu.edu	Tomasz Mrowka	Massachusetts Institute of Technology mrowka@math.mit.edu
Tobias H Colding	Massachusetts Institute of Technology colding@math.mit.edu	Aaron Naber	Northwestern University anaber@math.northwestern.edu
Simon Donaldson	Imperial College, London s.donaldson@ic.ac.uk	Peter Ozsváth	Princeton University petero@math.princeton.edu
Yasha Eliashberg	Stanford University eliash-gt@math.stanford.edu	Leonid Polterovich	Tel Aviv University polterov@post.tau.ac.il
Benson Farb	University of Chicago farb@math.uchicago.edu	Colin Rourke	University of Warwick gt@maths.warwick.ac.uk
David M Fisher	Rice University davidfisher@rice.edu	Roman Sauer	Karlsruhe Institute of Technology roman.sauer@kit.edu
Mike Freedman	Microsoft Research michaelf@microsoft.com	Stefan Schwede	Universität Bonn schwede@math.uni-bonn.de
David Gabai	Princeton University gabai@princeton.edu	Natasa Sesum	Rutgers University natasas@math.rutgers.edu
Stavros Garoufalidis	Southern U. of Sci. and Tech., China stavros@mpim-bonn.mpg.de	Gang Tian	Massachusetts Institute of Technology tian@math.mit.edu
Cameron Gordon	University of Texas gordon@math.utexas.edu	Ulrike Tillmann	Oxford University tillmann@maths.ox.ac.uk
Jesper Grodal	University of Copenhagen jg@math.ku.dk	Nathalie Wahl	University of Copenhagen wahl@math.ku.dk
Misha Gromov	IHÉS and NYU, Courant Institute gromov@ihes.fr	Anna Wienhard	Universität Heidelberg wienhard@mathi.uni-heidelberg.de

See inside back cover or msp.org/gt for submission instructions.

The subscription price for 2024 is US \$805/year for the electronic version, and \$1135/year (+\$70, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Geometry & Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Geometry & Topology (ISSN 1465-3060 printed, 1364-0380 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

GT peer review and production are managed by EditFlow® from MSP.

PUBLISHED BY

 **mathematical sciences publishers**
nonprofit scientific publishing
<http://msp.org/>

© 2024 Mathematical Sciences Publishers

GEOMETRY & TOPOLOGY

Volume 28 Issue 5 (pages 1995–2482) 2024

- Shear-shape cocycles for measured laminations and ergodic theory of the earthquake flow 1995
AARON CALDERON and JAMES FARRE
- The asymmetry of Thurston’s earthquake flow 2125
FRANCISCO ARANA-HERRERA and ALEX WRIGHT
- The persistence of a relative Rabinowitz–Floer complex 2145
GEORGIOS DIMITROGLOU RIZELL and MICHAEL G SULLIVAN
- Packing Lagrangian tori 2207
RICHARD HIND and ELY KERMAN
- The parabolic Verlinde formula: iterated residues and wall-crossings 2259
ANDRÁS SZENES and OLGA TRAPEZNIKOVA
- The signature and cusp geometry of hyperbolic knots 2313
ALEX DAVIES, ANDRÁS JUHÁSZ, MARC LACKENBY and NENAD TOMAŠEV
- Rigidity and geometricity for surface group actions on the circle 2345
KATHRYN MANN and MAXIME WOLFF
- Embedding surfaces in 4–manifolds 2399
DANIEL KASPROWSKI, MARK POWELL, ARUNIMA RAY and PETER TEICHNER