



Winter Braids Lecture Notes

Rob Schneiderman

Introduction to Whitney towers

Vol. 7 (2020), Course n° IV, p. 1-71.

<https://doi.org/10.5802/wbln.36>




Texte mis en ligne dans le cadre du
Centre Mersenne pour l'édition scientifique ouverte
<http://www.centre-mersenne.org/>
e-ISSN : 2426-0312

Introduction to Whitney towers

ROB SCHNEIDERMAN

Abstract

These introductory notes on Whitney towers in 4-manifolds, as developed in collaboration with Jim Conant and Peter Teichner, are an expansion of three expository lectures given at the Winter Braids X conference February 2020 in Pisa, Italy. Topics presented include local manipulations of surfaces in 4-space, fundamental definitions related to Whitney towers and their associated trees, geometric Jacobi identities, the classification of order n twisted Whitney towers in the 4-ball and higher-order Arf invariants, and low-order Whitney towers on 2-spheres in 4-manifolds and related invariants.



Introduction

These notes are an expansion of three introductory lectures given at the Winter Braids X conference February 2020 in Pisa, Italy describing a theory of Whitney towers on immersed surfaces in 4-manifolds, as developed in collaboration with Jim Conant and Peter Teichner. A Whitney tower is built on an immersed surface by iteratively adding Whitney disks pairing intersection points among the “higher-order” layers of Whitney disks and the surface (see section 1.8). The theory’s main goal is to use Whitney towers to gain topological information about the underlying immersed surface. These notes are focussed on providing a detailed introduction to the theory while simultaneously describing some open problems.

Section 1 covers local manipulations of surfaces and Whitney disks in 4-space, and fundamental definitions related to Whitney towers, including the associated trivalent trees that organize Whitney towers. This section culminates with the description of a geometric Jacobi Identity in the setting of Whitney towers.

Section 2 describes the classification of order n twisted Whitney towers on properly immersed disks in the 4-ball, which illustrates the order-raising intersection-obstruction theory and leads to the formulation of open problems related to certain higher-order Arf invariants which are invariants of classical link concordance. Here the trees associated to Whitney towers are seen to represent invariants in abelian groups related to Milnor's classical link invariants.

Section 3 reviews the classical intersection and self-intersection (order 0) homotopy invariants of 2-spheres in a 4-manifold X before introducing order 1 generalizations of these invariants in the setting of Whitney towers. Here new subtleties coming from $\pi_1 X$ and $\pi_2 X$ enter the picture. The end of this section describes open problems on the realization of the order 1 invariants when X is closed and $\pi_1 X$ is non-trivial.

The appendix section 4 provides additional material related to Section 2 and Section 3 in the form of outlines and/or details of proofs of results from those sections.

Exercises appear at the end of each section.

Sections 2 and 3 are largely independent of each other, but both depend on Section 1.

Conventions: Manifolds and submanifolds are assumed to be smooth, with generic intersections, unless otherwise specified, and during cut-and-paste constructions corners will be assumed to be rounded. The discussion throughout will also hold in the flat topological category via the notions of 4-dimensional topological transversality from [13, chap.9]. Orientations will usually be assumed but suppressed unless needed.

Acknowledgments: The author is supported by a Simons Foundation *Collaboration Grant for Mathematicians*. Also thanks to the organizers of the *Winter Braids X* conference, and of course collaborators Jim Conant and Peter Teichner.

Contents

Introduction	1
1. Disks in the 4-ball	3
1.1. Local coordinates	4
1.2. Signs of intersections	5
1.3. Finger moves	5
1.4. Whitney disks	5
1.5. Framed and twisted Whitney disks.	6
1.6. Parallel Whitney circles and disks	7
1.7. Whitney moves	7
1.8. Whitney towers	9
1.9. Trees	11
1.10. Trees for Whitney disks and intersection points	11
1.11. Twisted trees for twisted Whitney disks	12
1.12. 'Framed tree' terminology	13
1.13. Intersection forests of Whitney towers	13
1.14. Examples	13
1.15. Higher-order Whitney disks and intersections	15
1.16. Order n framed Whitney towers	15
1.17. Order n twisted Whitney towers	16
1.18. Other gradings of Whitney towers	16
1.19. Tree Orientations	17
1.20. 'Moving' unpaired intersections in their trees	18
1.21. Geometric Jacobi Identity in four dimensions	18
1.22. Section 1 Exercises	21
2. Order n twisted Whitney towers in the 4-ball	23
2.1. The order n twisted tree groups	24
2.2. Geometric meaning of the relations	25

2.3.	Intersection/obstruction theory for order n twisted Whitney towers	26
2.4.	Quick review of first non-vanishing Milnor invariants.	27
2.5.	The summation maps η_n	27
2.6.	The order n twisted Whitney tower filtration on links	28
2.7.	Definition of the realization maps	29
2.8.	Computing the order n twisted Whitney tower filtration	30
2.9.	The higher-order Arf invariant Conjecture	31
2.10.	Problems, Questions and re-formulations	32
2.11.	Higher-order Arf invariant Conjecture and Finite Type invariants	32
2.12.	Higher-order Arf invariant conjecture and transfinite Milnor invariants	33
2.13.	Section 2 Exercises	34
3.	Whitney towers on 2-spheres in 4-manifolds	36
3.1.	Classical intersection form	36
3.2.	Order 0 invariants	37
3.3.	Pulling apart pairs of spheres	38
3.4.	Decorated trees for order 1 intersections	39
3.5.	The order 1 non-repeating intersection invariant λ_1	39
3.6.	Pulling apart m -tuples of spheres	41
3.7.	The order 1 self-intersection invariant τ_1	42
3.8.	Examples and questions	43
3.9.	Section 3 Exercises	45
4.	Appendix	46
4.1.	Homotopy invariance of $\tau_1(A)$	46
4.2.	$\tau_1(A)$ vanishes if and only if A supports an order 2 framed Whitney tower	54
4.3.	Splitting twisted Whitney towers	58
4.4.	The Whitney move IHX relation	60
4.5.	The Whitney move twisted IHX relation	63
4.6.	Outline of twisted order-raising obstruction theory proof	65
4.7.	Whitney disk orientations and the AS relation	67
4.8.	Section 4 Exercises	68
	References	70

1. Disks in the 4-ball

This section focuses on local manipulations of surfaces in 4-space, and fundamental definitions related to Whitney towers, including the associated trivalent trees that organize Whitney towers. The section culminates with the description of a geometric Jacobi Identity in the setting of Whitney towers, followed by exercises related to the material covered here.

Throughout, the word “disk” means the 2-disk D^2 . A *sheet* of a surface in a 4-manifold is an embedded disk which is the properly embedded intersection of the surface with a 4-ball. Here a *proper* map sends boundary to boundary, and sends interior to interior. Following the tradition of Knot Theory, we will usually blur the distinction between a map and its image up to isotopy. In the case that the boundary ∂D of an immersed disk D is contained in the interior of an immersed surface A , we require and assume that ∂D is embedded, and also that the interior of D is disjoint from A near ∂D , i.e. that there exists a collar in D of ∂D such that the intersection of this collar with A is equal to ∂D . Orientations of surfaces and signs of intersection points will mostly be suppressed from notation.

1.1. Local coordinates

Figure 1.1 shows two properly embedded disks A and B in the 4-ball $B^4 = B^3 \times I$ which have a single transverse intersection point $p = A \pitchfork B$. In this figure the I -parameter increases from left to right, and each B^3 -slice $B^3 \times t$ for $t \in I$ intersects each disk in an arc. Thus $A \cup B$ is the track of a homotopy of arcs in B^3 which fixes the vertical black arc tracing out B , and isotopes the horizontal blue arc tracing out A from “front to back”, realizing a crossing change at $p = A \pitchfork B$.




Figure 1.1. Disks A and B in $B^4 = B^3 \times I$ with $p = A \pitchfork B$.

In the interest of avoiding visual clutter, we do not attempt to show the boundary of B^3 in any slice. So it is understood that the endpoints of the arcs lie in $\partial B^3 \times t$ for all t , since A and B are each properly embedded. Thus, to correctly interpret each picture in the movie one can either imagine the existence of the invisible 2-sphere $\partial B^3 \times t$ containing the visible endpoints of the arcs, or one can assume that the arcs continue further outside the picture until eventually reaching $\partial B^3 \times t$ which is lying outside the picture.

Thinking of the parameter t as “time”, we say that in Figure 1.1 the transverse intersection $p = A \pitchfork B$ is contained in the *present* slice $B^3 \times 0 \subset B^3 \times I$, and that both disks extend as arcs into *past* $B^3 \times -\epsilon$ and *future* $B^3 \times +\epsilon$. When describing surface sheets in this way we usually do not specify the endpoints of the interval I , which can be reparametrized as desired.




Figure 1.2. A and B in $B^4 = B^3 \times I$ with $p = A \pitchfork B$ and $A \subset B^3 \times 0$.

Figure 1.2 shows a different view of the same 4-ball neighborhood of $p = A \pitchfork B$ by changing coordinates and/or isotopy. In this figure the disk A is now contained in the present $B^3 \times 0 \subset B^3 \times I$, while B is still described by a fixed vertical arc times I . This view has the advantage of being completely described by just the present slice.

Again, A and B are each properly embedded in B^4 . So the rectangular appearance of A is either understood to extend further out of the picture until reaching the boundary 2-sphere of $B^3 \times 0$, or we can just take each $B^3 \times t$ to be a solid cube (with invisible PL boundary sphere).

Note that the transverse intersection p appears as a small “puncture” in the sheet A . This is just a visual device to help place the location of p in A . Also, the A -sheet is shown here as translucent, with a sub-arc of B visible “behind” A . In general, sheets in the present may appear translucent or opaque depending on which view provides a better clarification of the given configuration of sheets.

1.2. Signs of intersections

The sign $\epsilon_p \in \{+, -\}$ of an intersection p between oriented sheets A and B in an oriented 4-manifold X is defined to be $+$ (respectively, $-$) if the orientation of X at p agrees (respectively, disagrees) with the concatenation of the orientations of A and B at p .

1.3. Finger moves

A *finger move* is a regular homotopy supported near a *guiding arc* that creates a pair of oppositely-signed intersections between two sheets. This is illustrated in Figure 1.3 which shows a “movie of movies” starting from the disjoint sheets in the bottom and ending with the intersecting sheets in the top row. Each row shows the same 4-ball, with the bottom row showing the disjoint sheets before the finger move. Vertically and horizontally adjacent pictures are understood to be connected by smooth interpolations of the black arcs, with the vertical parameter corresponding to the homotopy, and the horizontal parameter corresponding to the I -parameter of $B^4 = B^3 \times I$ as in the previous figures. Between the middle row and the top row a single non-transverse tangential intersection will occur in the center picture at one moment of the homotopy.




Figure 1.3. A finger move guided by the red dashed arc.

Usually only the center top and/or center bottom pictures as in Figure 1.4 will be shown, with the other nearby and interpolating pictures shown in Figure 1.3 understood. Note that p and q have opposite signs for any choice of orientations on A , B and the ambient 4-ball (Exercise 1.22.4).

A finger move is supported in a 4-ball neighborhood of its guiding arc, so this picture also describes a finger move on two sheets in an arbitrary 4-manifold. Since this neighborhood can be taken to be arbitrarily close to the guiding arc, it may be assumed that the support of a finger move is disjoint from any other surfaces, and hence only creates the pair of intersections. Up to isotopy, a finger move is symmetric in the two sheets (Exercise 1.22.3).

1.4. Whitney disks

Let p and q be oppositely-signed transverse intersections between connected properly immersed surfaces A and B in a 4-manifold X , with p and q joined by embedded interior arcs $a \subset A$ and $b \subset B$ which are disjoint from all other singularities in A and B . We allow the possibility that $A = B$, in which case the circle $a \cup b$ must be embedded and change sheets at p




Figure 1.4. Before (left) and after (right) a finger move: The center-bottom and center-top pictures from Figure 1.3.

and q . Any immersed disk W in the interior of X bounded by such a *Whitney circle* $a \cup b$ is a *Whitney disk* pairing p and q .

For any collection W_i of Whitney disks pairing oppositely-signed $p_i, q_i \in A \natural B$ we require that the union $\cup_i \partial W_i$ of Whitney circles is embedded. This can always be arranged by controlled isotopies of the W_i (Exercise 1.22.14).




Figure 1.5. Two views of a model Whitney disk W pairing $p, q = A \natural B$.

Figure 1.5 shows a *model* Whitney disk W on sheets A and B in a 4-ball. In general, a Whitney disk in a 4-manifold has a neighborhood obtained by introducing *plumbings* [13, 1.2] into the model Whitney disk. So a Whitney disk may have interior self-intersections and intersections with other surfaces, but has an embedded collar which is disjoint from all surfaces, except for its boundary arcs $a \subset A$ and $b \subset B$.

An embedded Whitney disk whose interior is disjoint from all surfaces is called a *clean* Whitney disk.

1.5. Framed and twisted Whitney disks.

Let W be a Whitney disk pairing intersections $p, q \in A \natural B$, with boundary $\partial W = a \cup b$, for embedded arcs $a \subset A$ and $b \subset B$. Denote by $\nu_X W|_{\partial W}$ the restriction to ∂W of the normal disk-bundle $\nu_X W$ of W in X . Since p and q have opposite signs, $\nu_X W|_{\partial W}$ admits a nowhere-vanishing *Whitney section* $\overline{\partial W}$ defined by taking vectors tangent to A over a , and extending over b by vectors which are normal to B , as shown in the left of Figure 1.6.

The right side of Figure 1.6 shows $\overline{\partial W}$ inside an embedding into 3-space of $\nu_X W|_{\partial W} \cong S^1 \times D^2$. Although this choice of embedding has $\overline{\partial W}$ corresponding to the 0-framing of $D^2 \times S^1 \subset \mathbb{R}^3$, the section of $\nu_X W|_{\partial W}$ determined (up to homotopy) by the canonical framing of $\nu_X W$ will in general differ by full twists relative to $\overline{\partial W}$. (If p and q had the same sign then there would have to be a half-twist in the sheets of A and B , so the (continuous) Whitney section $\overline{\partial W}$ could not exist.)

If $\overline{\partial W}$ extends to a nowhere-vanishing section \overline{W} of $\nu_X W$, then W is said to be *framed* (since a disk-bundle over a disk has a canonical framing, and a nowhere-vanishing normal section over an oriented surface in an oriented 4-manifold determines a framing up to homotopy). In general, the obstruction to extending $\overline{\partial W}$ to a nowhere vanishing section of $\nu_X W$ is the relative Euler number $\chi(\nu_X W, \overline{\partial W}) \in \mathbb{Z}$, called the *twisting* of W and denoted $\omega(W)$, so W is




Figure 1.6. Left: Near a Whitney disk W pairing $p, q \in A \pitchfork B$, with $\partial W = a \cup b$, a Whitney section $\overline{\partial W}$ is shown in blue. This picture is accurate near ∂W ; in general, the interior of the evident (but not explicitly indicated) embedded W bounded by $a \cup b$ may have self-intersections as well as intersections with other surfaces. The line segments transverse to a in A indicate the correspondence with the right-hand picture of the normal disk-bundle $\nu_X W|_{\partial W}$. Right: The blue Whitney section $\overline{\partial W}$ is shown inside an embedding into 3-space of $\nu_X W|_{\partial W} \cong S^1 \times D^2$, with the sheets A and B indicated by line segments transverse to ∂W . The A -sheet cuts the front solid torus horizontally, while the B -sheet cuts the back of the solid torus vertically. The line segments transverse to b in B extend into past and future in the left-hand picture.

framed if and only if $\omega(W) = 0$. The twisting $\omega(W)$ can be computed by taking the intersection number of the zero section W with any extension \overline{W} of $\overline{\partial W}$ over W , so it does not depend on an orientation choice for W (since switching the orientation on W also switches the orientation of \overline{W}). And $\omega(W)$ is also unchanged by switching the roles of a and b in the construction of $\overline{\partial W}$, since interchanging the “tangent to...” and “normal to...” parts in the construction yields an isotopic section in $\nu_X W|_{\partial W}$ (isotopic through non-vanishing sections).

A Whitney disk W with $\omega(W) \neq 0$ is called a *twisted Whitney disk*.

1.6. Parallel Whitney circles and disks

The twisting $\omega(W)$ of W , which is the element of $\pi_1(\text{SO}(2)) \cong \mathbb{Z}$ determined by a Whitney section as described above, can also be computed using any section $\overline{\partial W}$ of $\nu_X W|_{\partial W}$ such that $\overline{\partial W}$ is in the complement of the tangent spaces of both A and B , since such $\overline{\partial W}$ will have the same number of rotations as a Whitney section (relative to the longitude determined by the canonical framing of ν_W); see Figure 1.6.

In these notes a *parallel Whitney circle* will refer to either these “nowhere tangent” sections, or the standard sections in section 1.5.

A disk is *parallel* to a Whitney disk if it is the extension of a parallel Whitney circle over the Whitney disk.

1.7. Whitney moves

Figure 1.7 describes a *model Whitney move*: In a 4-ball neighborhood of a framed embedded Whitney disk W , a pair of oppositely-signed transverse intersections p and q between surface sheets A and B is eliminated by a regular homotopy which isotopes one of the sheets across the clean framed W .

Combinatorially, the result of this Whitney move is constructed by deleting a regular neighborhood (in its sheet) of one arc of ∂W and replacing this neighborhood with a *Whitney bubble* over that arc. This Whitney bubble is formed from two parallel copies of W connected by a rectangular strip which is normal to a neighborhood (in its sheet) of the other arc. Figure 1.7 shows the result of adding a Whitney bubble to the sheet A . Although both these descriptions of the Whitney move involve a choice of arc of ∂W , up to isotopy the result is independent of this choice (Exercise 1.22.5).

A Whitney move on a clean framed W in a 4-manifold X is described by any embedding of this model into X which preserves the product structures and transversality of the sheets and W .




Figure 1.7. Left: W pairing $p, q = A \cap B$. Right: A Whitney move guided by W .

In general, a *Whitney move* on an arbitrary Whitney disk W is described using the combinatorial description of adding a Whitney bubble to one sheet: The Whitney bubble is still formed from two parallel copies of W together with a rectangular strip, and the intersections paired by W will still be eliminated by the Whitney move, but the following new intersections will be created. Each intersection $r \in W \cap C$ between the interior of W and a sheet C of another surface will give rise to two new intersections r' and r'' , as shown in Figure 1.8; each self-intersection of W will give rise to four new self-intersections; and if W is twisted, then $|\omega(W)|$ -many new self-intersections will be created corresponding to intersections between the two parallel copies of W in the Whitney bubble (see Exercise 1.22.7).




Figure 1.8

In our settings we will generally be able to arrange that Whitney disks are embedded and framed at the cost of creating interior intersections with sheets of other surfaces, so the Whitney move shown in Figure 1.8 will be the most relevant.

1.8. Whitney towers




Figure 1.9. Left: Four sheets and two Whitney disks in a 4-ball. Right: With the tree associated to the unpaired intersection.

A “successful” Whitney move eliminates the pair of intersections paired by the Whitney disk without creating any new intersections. Since this requires a clean framed Whitney disk, it is natural to try to somehow count interior intersections and twistings in an attempt to measure the obstructions to the existence of a homotopy of an immersed surface to an embedding.

Whenever these “higher-order intersections” can be paired by “higher-order Whitney disks” there is hope that they might be eliminated by successful “higher-order Whitney moves”, so one is led to construct a tower of Whitney disks by iteratively pairing up as many intersections as possible. As new Whitney disks are chosen, new intersections and twistings can appear which in turn need to be counted. We will see that invariants of the resulting Whitney tower can yield information about the bottom level immersed surface on which the tower is built.

As a basic example, the left side of Figure 1.9 shows four sheets in B^4 supporting a Whitney tower: The two intersections between red and blue are paired by a Whitney disk which itself has two interior intersections with green that are paired by another Whitney disk. This second Whitney disk has a single interior intersection with yellow that appears to be a “robust” obstruction to successfully homotoping these sheets rel boundary to be disjointly embedded in this 4-ball. In fact, we will see in Section 2 that this unpaired higher-order intersection does indeed represent such an obstruction (Exercise 2.13.7).

We can now state the general definition of a Whitney tower:

Definition 1.1. A Whitney tower on $A^2 \looparrowright X^4$ is defined by:

- (i) A itself is a Whitney tower.
- (ii) If \mathcal{W} is a Whitney tower and W is a Whitney disk pairing intersections in \mathcal{W} , then the union $\mathcal{W} \cup W$ is a Whitney tower.

In a Whitney tower the Whitney disk boundaries are required to be pairwise disjointly embedded (section 1.4).

Having constructed a Whitney tower by adding some number of Whitney disks as per the definition, we call the properly immersed surface A the *underlying surface* of the Whitney tower, and say that A *supports* the Whitney tower. The “boundary of \mathcal{W} ” is the boundary of A . We assume that A comes with a given orientation, and that \mathcal{W} is *oriented* by choosing (arbitrarily) and fixing orientations of all the Whitney disks in \mathcal{W} . Specific orientations and

signs will generally be suppressed from notation and discussion until needed, and at a first reading can be safely ignored while absorbing the combinatorics of Whitney towers.

Our underlying surfaces A will always be collections of disks or annuli or spheres, with components A_i indexed by $i \in \{1, 2, \dots, m\}$.

Definition 1.1 describes a very general notion of Whitney tower, and all of the figures so far give very basic examples that satisfy this definition. Structure will emerge upon the introduction of the univalent trees that organize Whitney towers (Figure 1.9 right, and Figure 1.10 bottom-right), and we will see that Whitney towers are interesting in their own right, as well as providing information about their underlying surfaces.




Figure 1.10. Zig-zagging down from top-left to bottom-left: Splitting a Whitney tower by finger moves. Bottom-right: The trees (blue) associated to un-paired intersections.

As motivation for introducing formalism to organize Whitney towers, Figure 1.10 shows how after “splitting” a Whitney tower by finger moves it can be arranged that all of its singularities are contained in regular neighborhoods of univalent trees, with each Whitney disk containing only one “problem” (un-paired intersection or Whitney disk boundary-arc). In section 4.3 this notion of splitting is described precisely and extended to include splitting of twisted Whitney disks.

1.9. Trees

A tree will always refer to a finite oriented univalent tree, where the orientation of a tree is given by cyclic orderings of the adjacent edges around each trivalent vertex. Univalent vertices will usually be labeled from the index set $\{1, 2, 3, \dots, m\}$, which will always correspond to the connected components of a properly immersed surface $A = A_1 \cup A_2 \cup A_3 \cup \dots \cup A_m \looparrowright (X, \partial X)$ supporting a Whitney tower. Trees are considered up to isomorphisms which preserve labels and orientations.

We identify formal non-associative iterated bracketings of elements from the index set $\{1, 2, 3, \dots, m\}$ with *rooted trees*, which have each univalent vertex labeled by an element from the index set, except for one *root* univalent vertex which is left unlabeled. So for instance the bracket (i, j) denotes the rooted tree $\begin{matrix} i \\ \diagup \\ j \end{matrix}$, for $i, j \in \{1, 2, 3, \dots, m\}$; and inductively (I, J) denotes $\begin{matrix} I \\ \diagup \\ J \end{matrix}$, where I and J denote bracketings/rooted trees. This is formalized in the first part of the following definition:

Definition 1.2. Let I and J be two rooted trees.

- (i) The *rooted product* (I, J) is the rooted tree gotten by identifying the root vertices of I and J to a single vertex v and sprouting a new rooted edge at v . This operation corresponds to the formal bracket (Figure 1.11 upper right). The orientation of (I, J) is inherited from those of I and J , and at v by the (left-to-right) ordering of the bracket product.
- (ii) The *inner product* $\langle I, J \rangle$ is the (unrooted) tree gotten by identifying the roots of I and J to a single non-vertex point. Note that $\langle I, J \rangle$ inherits an orientation from I and J , and that all the univalent vertices of $\langle I, J \rangle$ are labeled. (Figure 1.11 lower right.)




Figure 1.11. The *rooted product* (I, J) and *inner product* $\langle I, J \rangle$ of $I = (I_1, I_2)$ and $J = (J_1, J_2)$. All trivalent orientations correspond to a counterclockwise orientation of the plane.

1.10. Trees for Whitney disks and intersection points

Let $A = A_1, \dots, A_m \looparrowright (X, \partial X)$ be a properly immersed surface supporting a Whitney tower \mathcal{W} (oriented), where the A_i are the connected components of A .

To each component A_i is associated the rooted tree i consisting of a single edge with one vertex labeled by i , and to each transverse intersection $p \in A_i \pitchfork A_j$ is associated the tree $t_p := \langle i, j \rangle = i - j$ consisting of an edge with vertices labeled by i and j . Note that for singleton brackets (rooted edges) we drop the bracket from notation, writing i for (i) .

The rooted Y-tree $(i, j) = \begin{matrix} i \\ \diagup \\ j \end{matrix}$, with a single trivalent vertex and two univalent labels i and j , is associated to any Whitney disk $W_{(i,j)}$ pairing intersections between A_i and A_j . This rooted tree can be thought of as being embedded in \mathcal{W} , with its trivalent vertex and rooted edge sitting in $W_{(i,j)}$, and its two other edges descending into A_i and A_j as sheet-changing paths

(Figure 1.12 left). The cyclic orientation at the trivalent vertex corresponds to the orientation of $W_{(i,j)}$ via orientation conventions that will be described in section 1.19.

Note that Figure 1.12 shows the j -labeled univalent vertex of the tree (i, j) in the present, but this j -labeled edge changes sheets into A_j so the j -labeled univalent vertex is really in the past or future.




Figure 1.12. Left: The rooted tree $(i, j) \subset \mathcal{W}$ associated to $W_{(i,j)}$ (oriented counter-clockwise from the reader's point of view). Right: The unrooted tree $t_p = \langle (i, j), k \rangle \subset \mathcal{W}$ associated to $p \in W_{(i,j)} \pitchfork A_k$.

Associated to any transverse intersection $p \in W_{(i,j)} \pitchfork A_k$ is the unrooted tree

$$t_p = \langle (i, j), k \rangle = \overset{i}{\curvearrowright} \underset{j}{\curvearrowleft} \text{---} k,$$

as illustrated in Figure 1.12 right. We take the root vertex of each of the rooted trees associated to $W_{(i,j)}$ and A_k to be at p , so the inner product $t_p = \langle (i, j), k \rangle$ is realized geometrically as the union $\overset{i}{\curvearrowright} \underset{j}{\curvearrowleft} \cup_p \text{---} k$ of these rooted trees in \mathcal{W} .

Recursively, the rooted tree (I, J) is associated to any Whitney disk $W_{(I,J)}$ pairing intersections between W_I and W_J (see left-hand side of Figure 1.13); with the understanding that if, say, I is just a singleton i , then $W_I = W_i$ denotes the surface component A_i . And to any




Figure 1.13

transverse intersection $p \in W_{(I,J)} \cap W_K$ between $W_{(I,J)}$ and any W_K is associated the unrooted tree $t_p := \langle (I, J), K \rangle$ (see right-hand side of Figure 1.13).

The above description of the trees $t_p \subset \mathcal{W}$ can always be arranged to be disjointly embedded in \mathcal{W} , for any number of unpaired intersections in \mathcal{W} (Exercise 1.22.15).

1.11. Twisted trees for twisted Whitney disks

So far we have associated rooted trees to Whitney disks, and unrooted trees to unpaired intersection points. Recall from section 1.5 that a Whitney disk W_j with non-zero twisting

$\omega(W_j) \neq 0 \in \mathbb{Z}$ is said to be *twisted*. To keep track of twisted Whitney disks, we introduce one more type of tree:

For any rooted tree J , define the ∞ -tree

$$J^\infty := J \text{---} \infty$$

by labeling the root of J with the “twist” symbol ∞ . These ∞ -trees are called *twisted trees* since they are associated to twisted Whitney disks:

To each twisted Whitney disk $W_j \subset \mathcal{W}$ we associate the twisted tree J^∞ .

1.12. ‘Framed tree’ terminology

Since twisted trees are associated to twisted Whitney disks, for clarity we will sometimes refer to the unrooted trees t_p associated to intersection points as “framed” trees. In fact, section 4.3 shows that it can always be arranged that every unpaired intersection $p \in \mathcal{W}$ is an intersection between *framed* Whitney disks or surface components. (Also, “framed” is more succinct than “non- ∞ unrooted” or “not-twisted unrooted”.)

1.13. Intersection forests of Whitney towers

The above-described associations of trees capture the essential structure of Whitney towers, and the following key definition will be useful for both controlling constructions and defining invariants:

Definition 1.3. *The intersection forest $t(\mathcal{W})$ of a Whitney tower \mathcal{W} is the multiset of signed trees associated to unpaired intersections $p \in \mathcal{W}$ and \mathbb{Z} -coefficient ∞ -trees associated to twisted Whitney disks $W_j \subset \mathcal{W}$:*

$$t(\mathcal{W}) := \sum \epsilon_p \cdot t_p + \sum \omega(W_j) \cdot J^\infty$$

with $\epsilon_p \in \{+, -\}$ the usual sign of the transverse intersection point p , and $\omega(W_j) \in \mathbb{Z}$ the twisting of W_j .

Here the formal sums are over all unpaired p and all twisted W_j in \mathcal{W} . In fact, by splitting twisted Whitney disks (section 4.3) it can be arranged that all $\omega(W_j) = \pm 1$, just like the signs of the p .

In some papers $t(\mathcal{W})$ is notated as a disjoint union rather than a formal sum. Regardless of notation choice, it is helpful to think of $t(\mathcal{W})$ as being embedded in \mathcal{W} as per Figures 1.12 and 1.13 (and Exercise 1.22.15), since modifications of Whitney towers are described in terms of changes to $t(\mathcal{W})$, and these modifications correspond to relations in the targets of invariants represented by $t(\mathcal{W})$ that live in abelian groups generated by trees.

Remark 1.4 (Key Question). *Given a properly immersed surface A in a 4-manifold, constructing a Whitney tower \mathcal{W} on A involves many possible choices of Whitney disks, and since the essential structure of \mathcal{W} is contained in $t(\mathcal{W})$, it is natural to ask: “What does $t(\mathcal{W})$ tell us about A ?”. Various answers to this question have been found, as will be illustrated by results throughout these notes. More answers are waiting to be discovered...*

1.14. Examples

It will be shown in Lemma 2.11 that for any multiset of signed trees there exists a link $L \subset S^3$ bounding immersed disks into B^4 supporting a Whitney tower \mathcal{W} such that $t(\mathcal{W})$ is equal to the given multiset of signed trees.

On the other hand, Exercise 2.13.10 will show that there exist restrictions on the possible $t(\mathcal{W})$ for a Whitney tower \mathcal{W} on immersed 2-spheres in B^4 .

Some basic examples follow:




Figure 1.14

1.14.1.

Moving radially into B^4 from left to right, Figure 1.14 shows the Borromean Rings $L = L_1 \cup L_2 \cup L_3 \subset S^3 = \partial B^4$ bounding properly immersed disks $A = D_1 \cup D_2 \cup D_3 \looparrowright B^4$ which support $\mathcal{W} = D_1 \cup D_2 \cup D_3 \cup W_{(1,2)}$ with $t(\mathcal{W}) = \frac{1}{2} \succ 3$.

The disk D_1 consists of the “horizontal” opaque disk in the lower part of the middle picture extended by an annular collar back to L_1 in the left picture. The disks D_2 and D_3 consist of the embedded annuli which are the products of L_2 and L_3 with the radial coordinate into B^4 together with embedded disks (not shown) extending further into B^4 that are bounded by the unlink in the right picture. The framed embedded Whitney disk $W_{(1,2)}$ which pairs $D_1 \cap D_2$ is completely contained in the middle picture and has a single interior intersection point p with D_3 such that $t_p = \langle (1, 2), 3 \rangle = \frac{1}{2} \succ 3 = t(\mathcal{W})$.

1.14.2.

Moving radially into B^4 from left to right, Figure 1.15 shows the Figure-8 knot in S^3 bounding a properly immersed disk $A = D_1 \looparrowright B^4$ supporting $\mathcal{W} = D_1 \cup W_{(1,1)}$ with $t(\mathcal{W}) = (1, 1)^\infty = \frac{1}{1} \succ \infty$.

The track of the indicated null-homotopy of the Figure-8 knot describes the properly immersed disk D_1 with two self-intersections that are paired by the clean +1-twisted Whitney disk $W_{(1,1)}$. Part of $W_{(1,1)}$ is visible in the middle picture, and the unlink in the right hand picture can be capped off by two embedded disks which form the rest of D_1 and $W_{(1,1)}$. The twisting $\omega(W_{(1,1)}) = 1$ of $W_{(1,1)}$ is explained by Figure 1.16.




Figure 1.15




Figure 1.16. The two right-most pictures from Figure 1.15, but here with a Whitney section $\partial W_{(1,1)}$ shown in blue. Compare Figure 1.6. The +1-linking between the purple and blue circles corresponds to the twisting $\omega(W_{(1,1)}) = 1$.

1.14.3.

If \mathcal{W} is a Whitney tower on A such that $t(\mathcal{W}) = \emptyset$, then A is regularly homotopic to an embedding (Exercise 1.22.10).

1.15. Higher-order Whitney disks and intersections

We will be using the following grading of our trees:

Definition 1.5.

The order of a tree, rooted or unrooted, is defined to be the number of trivalent vertices.

Having associated trees to Whitney disks and intersection points in section 1.10, we can define *higher-order Whitney disks* and *higher-order intersections* in Whitney towers:

Definition 1.6.

The order of a Whitney disk or intersection point in a Whitney tower is defined to be the order of its corresponding tree.

The components A_i of the underlying surface are also referred to as *order 0 surfaces*.

Linking numbers of links in S^3 can be computed by summing order 0 intersections between order 0 surfaces in B^4 bounded by the links. The higher-order intersections and Whitney disk twistings in a Whitney tower bounded by a link are closely related to “higher-order” linking numbers (Milnor’s link invariants), as will be explained in Section 2.

1.16. Order n framed Whitney towers

Definition 1.7. \mathcal{W} is an order n framed Whitney tower if

- (i) every framed tree t_p in $t(\mathcal{W})$ is of order $\geq n$, and
- (ii) there are no ∞ -trees in $t(\mathcal{W})$.

So in an order n framed \mathcal{W} all unpaired intersections have order $\geq n$, and all Whitney disks are framed.

1.17. Order n twisted Whitney towers

Definition 1.8. \mathcal{W} is an order n twisted Whitney tower if

- (i) every framed tree t_p in $t(\mathcal{W})$ is of order $\geq n$, and
- (ii) every ∞ -tree in $t(\mathcal{W})$ is of order $\geq \frac{n}{2}$.

So in an order n twisted \mathcal{W} all unpaired intersections have order $\geq n$, and all Whitney disks of order less than $n/2$ are framed.

The reason that Definition 1.8 allows twisted Whitney disks in orders at least $n/2$ will become clear in Section 2, where it will be shown (following [8]) that intersection invariants extracted from the intersection forests of order n twisted Whitney towers, and the corresponding order-raising obstruction theory, lead to a classification of order n twisted Whitney towers on properly immersed disks in the 4-ball in terms of Milnor invariants and higher-order Arf invariants of the links on the boundary. In [8] an analogous classification of order n framed Whitney towers is also derived from this twisted classification.

1.18. Other gradings of Whitney towers

We mention here just two of the other variations on organizing Whitney towers. See also [8, Sec.1.5] for a brief comparison discussion.

1.18.1. Non-repeating order n Whitney towers:

If t_p is the tree associated to an intersection p in a Whitney tower such that t_p has distinct univalent labels, then t_p is called a *non-repeating tree* and p is called a *non-repeating intersection*.

Definition 1.9. A Whitney tower \mathcal{W} is an order n non-repeating Whitney tower if all non-repeating $t_p \in t(\mathcal{W})$ are of order $\geq n$.

Note that there are no restrictions on ∞ -trees in order n non-repeating Whitney towers, and no restrictions on any *repeating trees* which do not have distinctly labeled univalent vertices.

Non-repeating Whitney towers characterize being able to “pull apart” components:

Theorem 1.10 ([31]). $A = \cup_{i=1}^m A_i \looparrowright X$ supports an order $m - 1$ non-repeating \mathcal{W} if and only if A is homotopic rel ∂A to $A' = \cup_{i=1}^m A'_i$ with $A'_i \cap A'_j = \emptyset$ for all $i \neq j$.

The obstructions to the existence of order n non-repeating Whitney towers are analogous to Milnor’s non-repeating link homotopy invariants [22], and for immersed disks in the 4-ball are equivalent to the first non-vanishing non-repeating Milnor invariants [31, Thm.8].

In Section 3 we will examine the obstructions to non-repeating Whitney towers on 2-spheres in 4-manifolds in detail for low orders.

1.18.2. Symmetric Whitney towers:

Definition 1.11. A Whitney tower \mathcal{W} is symmetric if all Whitney disks in \mathcal{W} are framed and only have interior intersections with Whitney disks of the same order.

The natural grading of symmetric Whitney towers is by *height*:

Definition 1.12. A symmetric Whitney tower \mathcal{W} has height n if \mathcal{W} has order $(2n - 2)$ as a framed Whitney tower.

The Whitney disks in a symmetric Whitney tower correspond to symmetric rooted trees Y^n defined in Figure 1.17.



Figure 1.17. The symmetric rooted trees Y^1, Y^2, Y^3 and $Y^n = \langle Y^{(n-1)}, Y^{(n-1)} \rangle$ of height 1, 2, 3, and n .

So an equivalent definition of a symmetric Whitney tower of height n is that $t(\mathcal{W})$ only contains trees which are inner products $\langle Y^n, Y^n \rangle$ of height n symmetric rooted trees.

We also have the following variation extending this definition to half-integers: If $t(\mathcal{W})$ only contains trees which are inner products $\langle Y^n, Y^{n-1} \rangle$, then \mathcal{W} is said to be of height $n.5$.

Theorem 1.13 ([6]). *If a link $L \subset S^3$ bounds immersed disks into B^4 supporting \mathcal{W} of height $n + 2$, then L is n -solvable in the sense of Cochran–Orr–Teichner.*

A knot bounds a height 2 Whitney tower if and only if it has vanishing Arf invariant; a knot bounds a height 2.5 Whitney tower if only if it is algebraically slice; and in general signature invariants are known to provide upper bounds on the possible heights of Whitney towers bounded by knots [6]. There is an extensive literature related to the n -solvable filtration of classical concordance.

A complete height-raising obstruction theory for symmetric Whitney towers is not known. Necessary conditions for raising height can be formulated in terms of the vanishing of signatures, but algebraic invariants whose vanishing suffices for raising height have not been formulated in general. The order-raising construction given in section 4.6 destroys any Whitney disk symmetry. First open case: Given a knot in S^3 bounding $\mathcal{W} \subset B^4$ of height 2.5, there are no known obstructions to the knot bounding a Whitney tower of height 3 in B^4 .

1.19. Tree Orientations

Recall that our Whitney towers are oriented, with the underlying surface A coming with a given orientation, and with fixed orientations chosen for the Whitney disks in \mathcal{W} . With an eye towards later defining invariants in terms of $t(\mathcal{W})$ which will be independent of Whitney disk orientations, we introduce here two conventions from the Whitney tower literature which relate Whitney disk and tree orientations. Via either of these two conventions, the choices of Whitney disk orientations will be seen in section 4.7 to correspond to *antisymmetry relations* among the trees, and will be quotiented out in the target of the invariants.

As per section 1.10, the rooted tree associated to a Whitney disk W in \mathcal{W} can be mapped into \mathcal{W} , with the trivalent vertex adjacent to the root contained in the interior of W , and each other trivalent vertex contained in the interior of a lower-order Whitney disk supporting W . The two descending edges from each trivalent vertex determine a “corner” of the corresponding Whitney disk that does not contain the third ascending edge. For instance, in Figure 1.13 the corner determined by the descending edges from the indicated trivalent vertex contains the negative intersection point paired by the corresponding Whitney disk. And in the right side of Figure 1.9 the corners of the Whitney disks determined by the descending edges of the two trivalent vertices of the tree each contain the negative intersection paired by the Whitney disk.

With this terminology we define the following conventions for aligning the vertex orientations on trees with the chosen orientations on Whitney disks:

The negative corner convention: The trees associated to all Whitney disks in \mathcal{W} are mapped into \mathcal{W} with the requirement that the corner of each Whitney disk determined by the two descending edges of each trivalent vertex contains the *negative* intersection point

between the sheets paired by the corresponding Whitney disk. Then the vertex orientations of each tree are taken to be induced from the Whitney disk orientations.

The positive corner convention: The trees associated to all Whitney disks in \mathcal{W} are mapped into \mathcal{W} with the requirement that the corner of each Whitney disk determined by the two descending edges of each trivalent vertex contains the *positive* intersection point between the sheets paired by the corresponding Whitney disk. Then the vertex orientations of each tree are taken to be induced from the Whitney disk orientations.




Figure 1.18. Top: A $W_{(I,J)}$ -Whitney move on the I -sheet in a split Whitney tower creates a pair of intersections which can be paired by a ‘meridional’ Whitney disk $W_{(K,I)}$ that intersects the J -sheet in a single point. Bottom: Indicating the unpaired intersection by a small circle linking its corresponding edge, the top construction realizes $\langle (I,J), K \rangle = \langle (K,I), J \rangle$.

1.20. ‘Moving’ unpaired intersections in their trees

Again with an eye towards extracting invariants from $t(\mathcal{W})$, examination of our notation reveals that the tree $t_p := \langle (I,J), K \rangle$ associated to $p \in W_{(I,J)} \# W_K$ does not keep track of which edge of t_p contains the unpaired intersection p , since for instance $\langle (I,J), K \rangle = \langle I, (J,K) \rangle = \langle (K,I), J \rangle$. It seems reasonable ask “Why not also keep track of the edge in t_p corresponding to p ?”, but it turns out that robust information is carried by the shape of the tree rather than the intersection point itself. This is illustrated in Figure 1.18, which shows the geometric realization of $\langle (I,J), K \rangle = \langle (K,I), J \rangle$ by a controlled local construction that does not create any new trees. This construction preserves the signed tree of the unpaired intersection point, using either convention in section 1.19 (Exercise 1.22.11). By iterating this construction an un-paired intersection can be “moved to any edge of its tree”.

1.21. Geometric Jacobi Identity in four dimensions

We close this section by showing the necessity of certain relations in the target of any homotopy invariant of A represented by $t(\mathcal{W})$ for \mathcal{W} supported by A . The proof of the following theorem provides a nice illustration of the material covered so far:

Theorem 1.14 ([7]). *There exist four 2-spheres $A_1 \cup A_2 \cup A_3 \cup A_4 \looparrowright S^4$ in the 4-sphere supporting \mathcal{W} with intersection forest $t(\mathcal{W})$ equal to the three signed trees in Figure 1.19.*




Figure 1.19

This is a “Jacobi Identity” in the sense that taking the 4-labeled univalent vertices as roots and identifying the resulting rooted trees with Lie brackets yields the three terms $[[1, 2], 3] - [1, [2, 3]] + [[3, 1], 2]$, and this sum of trees must vanish in the target of any homotopy invariant represented by $t(\mathcal{W})$ since \mathcal{W} is supported by homotopically trivial 2-spheres.

Proof of Theorem 1.14. Start with disjoint embeddings $A_i : S^2 \rightarrow B^4 \subset S^4$, $i = 1, 2, 3, 4$.

Then do a single finger move on each of A_1, A_2, A_3 into A_4 , yielding the left picture of Figure 1.20, in which A_4 appears as the horizontal plane of each picture. The Whitney disks $W_{(3,4)}$, $W_{(4,1)}$ and $W_{(2,4)}$ in the center picture of Figure 1.20 are inverse to the finger moves (Exercise 1.22.6).




Figure 1.20

We will describe modifications of the order 1 Whitney disks in the center of Figure 1.20 that yield new disjointly embedded framed order 1 Whitney disks which will have the boundaries shown in the right picture of Figure 1.20. These modifications will create pairs of order 1 intersections that admit disjointly embedded framed order 2 Whitney disks which each have a single order 2 intersection corresponding to a term of the Jacobi identity. The entire construction will be supported in the 4-ball described by Figure 1.20 as the present, and extending into past and future where the arcs of A_1, A_2 and A_3 extend as products.

First we will construct the left-most term of the Jacobi identity in Figure 1.19, suppressing orientations for the moment. Start by changing a collar of $W_{(3,4)}$ as indicated in the left of Figure 1.21. Still referring to this new Whitney disk as $W_{(3,4)}$, we have created a new pair of order 1 intersections $\{q, r\} = A_2 \pitchfork W_{(3,4)}$, as shown in the figure. Since the collar of $W_{(3,4)}$ containing q and r is parallel to A_4 , it follows that q and r have opposite signs, since the intersections $A_2 \pitchfork A_4$ created by the finger move of A_2 into A_4 have opposite signs. Next add an order 2 Whitney disk $W_{(2,(3,4))}$ pairing q and r as on the right of Figure 1.21, where $\partial W_{(2,(3,4))}$ is shown as green. Part of $W_{(2,(3,4))}$ is formed from the original order 1 Whitney disk $W_{(2,4)}$ with a collar of the arc of $\partial W_{(2,4)}$ on A_4 removed, and the rest of $W_{(2,(3,4))}$ is parallel to A_4 . This creates a single order 2 intersection $p = A_1 \cap W_{(2,(3,4))}$, as shown in the figure. This new order 1 $W_{(3,4)}$ and the order 2 $W_{(2,(3,4))}$ are each embedded, and completely contained in the present.

In the left of Figure 1.22 the tree $t_p = \begin{matrix} 3 \\ \searrow \\ 4 \end{matrix} \succ \begin{matrix} 2 \\ \swarrow \\ 1 \end{matrix}$ associated to $p = A_1 \cap W_{(2,(3,4))}$ is included (but suppressed from view are the continuations of the edges changing sheets into the order 0 2-spheres). If $W_{(3,4)}$ and $W_{(2,(3,4))}$ are oriented to agree with the orientation of A_4 where they are parallel to A_4 , then this embedding of t_p conforms to the positive corner convention of

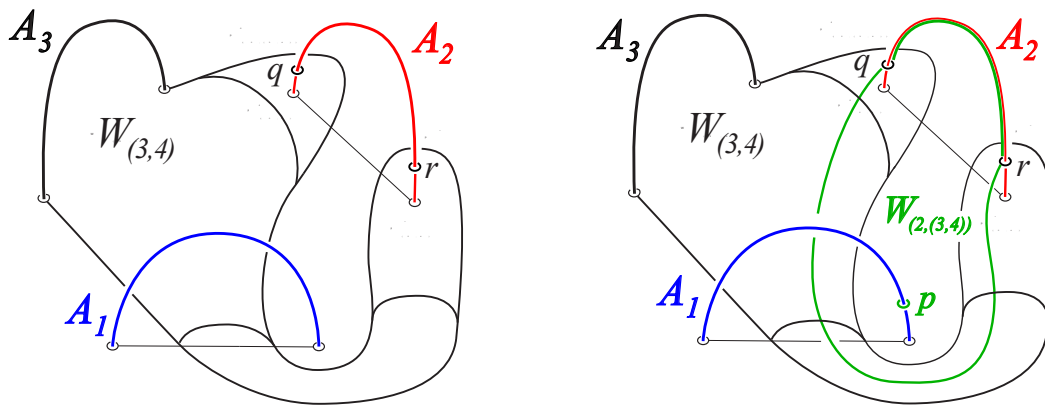


Figure 1.21

section 1.19, and the sign ϵ_p of p is positive. So $\epsilon_p \cdot t_p$ is the left-most term in the Jacobi relator of Figure 1.19.

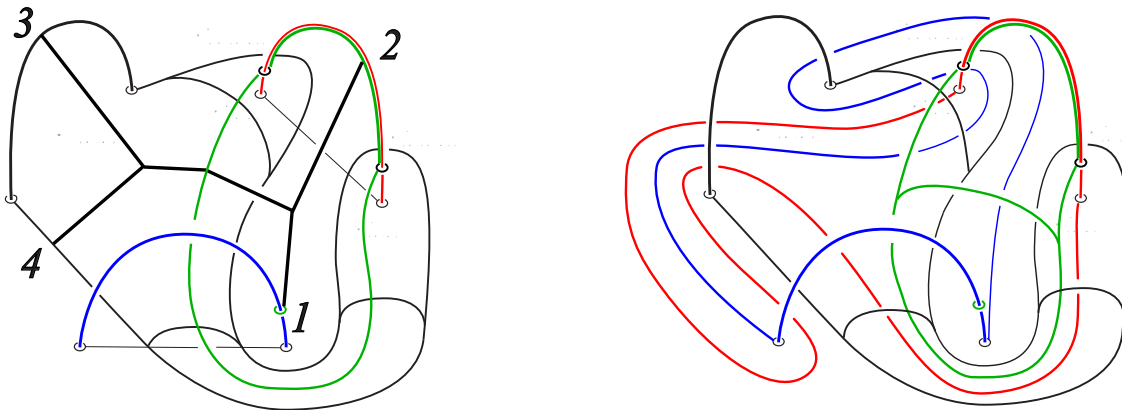


Figure 1.22

It remains to construct the other two terms in Figure 1.19. The right side of Figure 1.22 shows the just constructed $W_{(3,4)}$ and $W_{(2,(3,4))}$, and the red and blue boundaries of the new order 1 Whitney disks that we want to create. Observe that the parts of the red and blue boundaries that lie on A_4 extend to small embedded collars in the present that are disjoint from $W_{(3,4)}$ and $W_{(2,(3,4))}$ as well as the four 2-spheres. So by extending the inner collar boundary of red into the past, and the inner collar boundary of blue into future, the same construction that was just done in the present can be carried out in nearby past and future B^3 -slices to yield the other two desired trees. See Exercise 1.22.16.

Note that the construction necessarily extends into both past and future because it yields new disjointly embedded order 1 Whitney disks with boundaries as in the right of Figure 1.20, and this configuration forms the Borromean Rings which is not a slice link. \square

Corollary 1.15. *For any Whitney tower \mathcal{W} on a properly immersed surface A in a 4-manifold, the local 'IHX relation' of finite type theory (Figure 1.23) is needed in the target of any invariant of A represented by $t(\mathcal{W})$.*

In this more general local relation the univalent vertices of the three trees represent arbitrary (fixed) subtrees. To see that this corollary follows from Theorem 1.14, observe that clean Whitney disks W_I, W_J, W_K and W_L corresponding to any given rooted trees I, J, K and L can be created by finger moves (Exercise 1.22.12). Then, after a connected sum with S^4




Figure 1.23. The IHX relation.

along a 4-ball in the complement of \mathcal{W} , tubing the 2-spheres from the theorem into the interiors of these Whitney disks creates exactly the IHX relator in the corollary, since the tubes are supported near arcs which can be taken to be disjoint from the 2-complex \mathcal{W} .

1.22. Section 1 Exercises

1.22.1. Exercise:

Visualize the 2-component unlink in S^3 bounding the disjoint disks in B^4 shown in Figure 1.24.




Figure 1.24. Disjoint disks in $B^4 = B^3 \times I$

1.22.2. Exercise:

Visualize the Hopf link $= \partial A \cup \partial B \subset S^3 = \partial(B^3 \times I)$ in each of Figure 1.1 and Figure 1.2.

1.22.3. Exercise:

Starting with the bottom row of Figure 1.3, draw the finger move that pushes the horizontal blue sheet down into the black sheet. Observe that the results of the two finger moves are isotopic.

1.22.4. Exercise:

Check that the two intersections created by a finger move have opposite signs.

1.22.5. Exercise:

Check that Figure 1.25 shows the result of a model Whitney move which adds the Whitney bubble to the black sheet instead of the blue sheet in Figure 1.7. (The two copies of the Whitney disk in the Whitney bubble are in grey). Convince yourself that a model Whitney move is symmetric up to isotopy by finding an isotopy between Figure 1.25 and the right side of Figure 1.7.




Figure 1.25

1.22.6. Exercise:

The oppositely-signed pair of intersections created by a finger move (Figure 1.4) are contained in a local 4-ball and admit a model Whitney disk (Figure 1.5). Check that the result of doing a Whitney move on such a Whitney disk (Figure 1.7) is isotopic to not doing the finger move in the first place. We sometimes say that such a Whitney disk is “inverse” to the finger move.

1.22.7. Exercise:

Describe the new intersections created by a W -Whitney move in terms of the interior intersections W has with surfaces, the self-intersections of W , and the twisting $\omega(W)$.

1.22.8. Exercise:

Visualize the Borromean Rings $\partial A \cup \partial B \cup \partial C \subset S^3 = \partial(B^3 \times I)$ in (both sides of) Figure 1.8. HINT: See the Bing-double of one component of Exercise 1.22.2.

1.22.9. Exercise:

Visualize the Bing-double of the Hopf link as the boundaries of the red, blue, green and yellow sheets in $S^3 = \partial(B^3 \times I)$ in the left side of Figure 1.9.

1.22.10. Exercise:

Show that if \mathcal{W} is a Whitney tower on A such that $t(\mathcal{W}) = \emptyset$, then A is regularly homotopic to an embedding.

1.22.11. Exercise:

Check that the construction of Figure 1.18 preserves the sign of the unpaired intersection point using either the positive or negative convention in section 1.19. You may assume that the embedded trees in the figure satisfy the orientation convention, and check that this implies that the sign of the unpaired intersections are the same on the left and right.

1.22.12. Exercise:

Let \mathcal{W} be a Whitney tower on a surface A , and let I be any rooted tree. Show that, without changing $t(\mathcal{W})$, a clean framed Whitney disk W_I with associated tree I can be created by performing finger moves. (HINT: A finger move is supported near an arc, hence can be arranged to miss any other surface.)

1.22.13. Exercise:

Figure 1.26 shows how an intersection between a green sheet and the interior of a Whitney disk W can be eliminated by a finger move, at the cost of creating an oppositely-signed pair of intersections between the green sheet and one of the sheets paired by W . This is called “pushing down” an intersection in W . Use pushing down to show that if A supports an order 1 framed Whitney tower \mathcal{W} , then A supports an order 1 framed Whitney tower \mathcal{W}' whose Whitney disks are disjointly embedded.




Figure 1.26. ‘Pushing down’ an intersection.

1.22.14. Exercise:

The green-blue intersections created by the finger move in Figure 1.26 admit a local Whitney disk V which is inverse to the finger move (Exercise 1.22.6), such that ∂V intersects ∂W transversely in the blue sheet. Draw V into the right side of Figure 1.26 (hanging down “underneath” the blue sheet). Observe that V and W can be made disjoint by extending to a collar of V an isotopy of ∂V in the blue sheet that pushes $\partial V \pitchfork \partial W$ away across either blue-black intersection paired by W . This isotopy of V creates an interior intersection between (the new) V and the black sheet. Draw this new V , including its interior intersection with the black sheet, and check that this intersection between V and black has the same order as the original intersection between W and the green sheet.

1.22.15. Exercise:

Check that the trees $t_p \subset \mathcal{W}$ of section 1.10 can always be arranged to be disjointly embedded in \mathcal{W} , for any number of unpaired intersections in \mathcal{W} . (Do not assume that \mathcal{W} is split (Figure 1.10), and don’t forget the possibility of self-intersections in a Whitney disk.)

1.22.16. Exercise:

Complete the proof of Theorem 1.14 by constructing the other two trees of Figure 1.19 analogously using past and future (see [7] for the solution), and check that all the Whitney disks are framed.

2. Order n twisted Whitney towers in the 4-ball

Recall from Definition 1.8 that a Whitney tower \mathcal{W} is *twisted of order n* if every unpaired intersection p in \mathcal{W} is of order $\geq n$ and every twisted Whitney disk W_j in \mathcal{W} is of order $\geq \frac{n}{2}$; or equivalently, if every framed tree t_p in the intersection forest $t(\mathcal{W})$ is of order $\geq n$ and every twisted tree J^∞ in $t(\mathcal{W})$ is of order $\geq \frac{n}{2}$.

We say that a link $L \subset S^3 = \partial B^4$ bounds an order n twisted Whitney tower \mathcal{W} if $\mathcal{W} \subset B^4$ is an order n twisted Whitney tower whose order zero surfaces are immersed disks bounded by the components of L .

Following [8, 11], using algebraic results from [9], this section describes a classification of links in S^3 bounding order n twisted Whitney towers:

Section 2.1 defines abelian groups \mathcal{T}_n^∞ generated by trees, with relations corresponding to controlled modifications of twisted Whitney towers. These groups are the targets for intersection invariants $\tau_n^\infty(\mathcal{W}) := [t(\mathcal{W})] \in \mathcal{T}_n^\infty$ defined in section 2.3 which have the property that L bounds an order n twisted \mathcal{W} with $\tau_n^\infty(\mathcal{W}) = 0$ if and only if L bounds an order $n + 1$ twisted Whitney tower. The identification of these order-raising obstruction-theoretic invariants with Milnor invariants and higher-order Arf invariants of the link on the boundary leads to a classification of order n twisted Whitney towers in B^4 (Corollary 2.18).

The higher-order Arf invariants appear for each $n \equiv 2 \pmod{4}$, and take values in specific finite $\mathbb{Z}/2\mathbb{Z}$ -vector spaces (Definition 2.17). In this setting the main open problem is to determine precisely the image of these higher-order Arf invariants, which the *Higher-order Arf invariant Conjecture* states is maximal (Conjecture 2.19).

For links bounding order n framed Whitney towers (Definition 1.7) there is an analogous classification that is more complicated to describe because, in addition to Milnor invariants and higher-order Arf invariants, it also involves *higher-order Sato–Levine invariants* which represent obstructions to framing odd order twisted Whitney towers and correspond to certain projections of Milnor invariants. In [8] the framed classification is derived from the twisted one described here.

We also remark that Jae Choon Cha has shown that the higher-order Arf invariants measure the potential difference between the existence of twisted \mathcal{W} in B^4 versus rational homology B^4 s. Namely, $L \subset S^3$ has vanishing Milnor invariants through order n if and only if L bounds an order $n + 1$ twisted \mathcal{W} in a *rational homology* B^4 (see [3, Thm.C]).

Throughout this section the fixed index set $\{1, 2, \dots, m\}$ is usually suppressed from notation.

2.1. The order n twisted tree groups

Recall our terminology and conventions for trees (sections 1.9–1.13, and 1.15), including that *order* is the number of trivalent vertices.

First we define framed tree groups:

Definition 2.1. Denote by \mathcal{T}_n the free abelian group on order n framed trees modulo the local antisymmetry (AS) and Jacobi identity (IHX) relations in Figure 2.1.




Figure 2.1. The AS (left) and IHX (right) local relations. Here univalent vertices represent arbitrary fixed subtrees.

The target twisted tree groups \mathcal{T}_n^∞ for the intersection invariants $\tau_n^\infty(\mathcal{W})$ for order n twisted Whitney towers $\mathcal{W} \subset B^4$ bounded by $L \subset S^3$ will be defined separately for odd and even n . After giving the definitions in terms of generating trees and relations, the geometric meaning of the relations will be discussed.

Definition 2.2. For each $j \geq 1$, the order $2j - 1$ twisted tree group $\mathcal{T}_{2j-1}^\infty$ is the quotient of \mathcal{T}_{2j-1} by boundary-twist relations:

$$i \prec_j^j = 0$$

where J ranges over all order $j - 1$ subtrees.

Definition 2.3. For each $j \geq 0$, the order $2j$ twisted tree group \mathcal{T}_{2j}^∞ is the quotient of the free abelian group on framed trees of order $2j$ and ∞ -trees of order j by the following relations:

- (i) AS and IHX relations on order $2j$ framed trees
- (ii) symmetry relations: $(-J)^\infty = J^\infty$
- (iii) twisted IHX relations: $I^\infty = H^\infty + X^\infty - \langle H, X \rangle$
- (iv) interior-twist relations: $2 \cdot J^\infty = \langle J, J \rangle$

In item (iii) the three twisted trees differ locally as in the right of Figure 2.1.

See [10, Sec.4.8] for an interpretation of $J \mapsto J^\infty$ as a quadratic refinement of the \mathcal{T}_n -valued intersection form on rooted trees (or on Whitney disks).

2.2. Geometric meaning of the relations

Both the odd and even order twisted tree groups contain the AS and IHX relations which apply to framed tree generators. We have already seen in Corollary 1.15 the necessity of including the IHX relations in defining an invariant from $t(\mathcal{W})$ since IHX trees can be created locally. Upon fixing the positive or negative corner convention (section 1.19), the signs $\epsilon_p = \pm$ of the framed trees $\epsilon_p \cdot t_p$ in $t(\mathcal{W})$ only depend on the orientation of the underlying order 0 surface modulo the antisymmetry relations (section 4.7).

In the odd order groups $\mathcal{T}_{2j-1}^\infty$, which contain the obstructions to the existence of an order $2j$ twisted Whitney tower, the *boundary-twist relations* correspond geometrically to the fact that performing a boundary twist (Figure 2.2) on an order j Whitney disk $W_{(i,j)}$ creates an order $2j - 1$ intersection point $p \in W_{(i,j)} \cap W_j$ with associated tree $t_p = i \leftarrow_j$ and changes $\omega(W_{(i,j)})$ by ± 1 . By Exercise 1.22.12, any number of clean framed $W_{(i,j)}$ can be created in any Whitney tower, so any number of $t_p = i \leftarrow_j$ can be created by this construction. It follows from the obstruction theory (section 4.6) that after arranging such trees into “algebraically canceling” pairs, the corresponding unpaired intersections can be exchanged for “geometrically canceling” intersections admitting Whitney disks. Since this can be done at the cost of only creating order j twisted Whitney disks, which are allowed in an order $2j$ Whitney tower, the trees $i \leftarrow_j$ do not represent obstructions.

In the even order target groups \mathcal{T}_{2j}^∞ : The *symmetry relation* corresponds to the fact that the twisting $\omega(W)$ (section 1.5) is independent of the orientation of the Whitney disk W , with the minus sign denoting that the cyclic orderings at the trivalent vertices of $-J$ differ from those of J at an odd number of vertices. The *twisted IHX relation* corresponds to the effect of performing a Whitney move in the presence of a twisted Whitney disk, as described in Lemma 4.5 and [8, Lem.4.1]. The *interior-twist relation* corresponds to the fact that creating a ± 1 self-intersection in a W_j by a local cusp-homotopy [13, Sec.1.6] changes the twisting by ∓ 2 (Figure 2.3). The result of such a cusp-homotopy is the same as the local cut-and-paste operation described in [13, Sec.1.3]. For any J , a clean W_j can be created by finger moves (Exercise 1.22.12), then a \pm -interior twist W_j will change $t(\mathcal{W})$ by:

$$\pm \langle J, J \rangle \mp 2 \cdot J^\infty$$

Thus, all the relations in \mathcal{T}_n^∞ can be realized by controlled modifications of Whitney towers altering their intersection forests, without changing the homotopy class of the underlying order 0 surface.

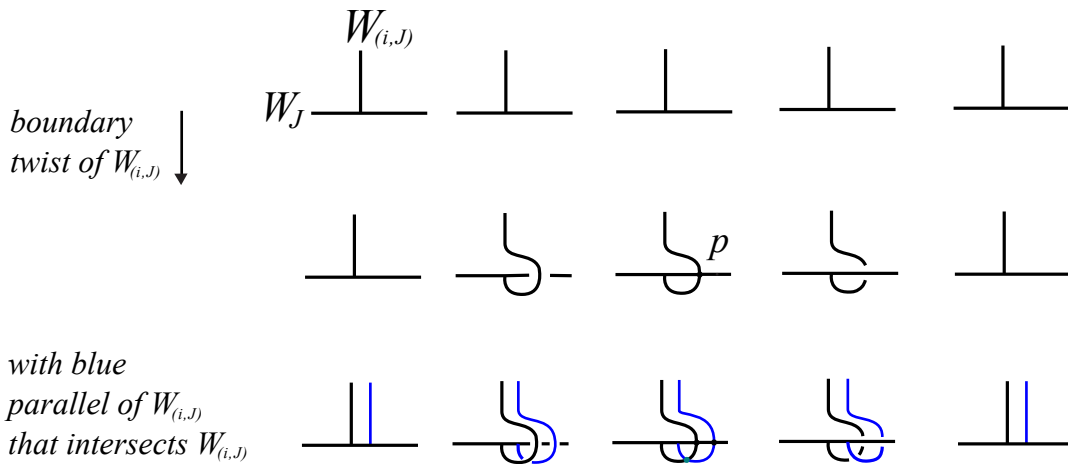


Figure 2.2. Top two rows: ‘Side view’ near a point in $\partial W_{(i,J)} \subset W_J$ of the boundary-twist operation on $W_{(i,J)}$ into W_J which creates a transverse intersection $p \in W_{(i,J)} \cap W_J$. Bottom row: This operation changes $\omega(W_{(i,J)})$ by ± 1 , as represented by the intersection in the middle frame between the boundary-twisted $W_{(i,J)}$ and its blue parallel copy.

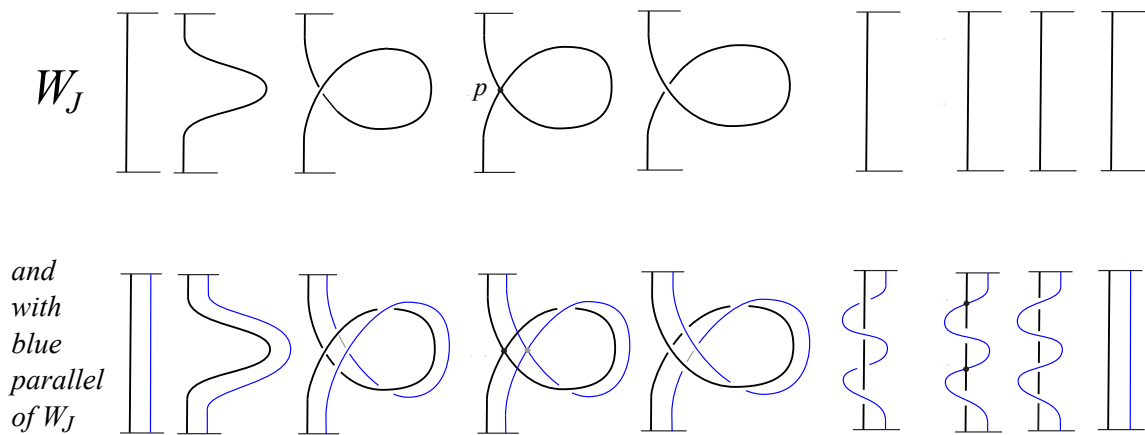


Figure 2.3. After a \pm -interior twist on W_J , shown near an arc in W_J that runs between the two sheets paired by W_J . Top row: A new transverse self-intersection $p \in W_J \pitchfork W_J$ has been created. Bottom row: $\omega(W_J)$ has been changed by ∓ 2 , as represented by the pair of blue-black intersections between this new W_J and a parallel copy which are visible in the third-from-right frame. Note that the pair of blue-black intersections in the fourth-from-left frame are just an artifact of the immersion of the normal bundle of W_J and do not contribute to the twisting $\omega(W_J)$.

2.3. Intersection/obstruction theory for order n twisted Whitney towers

Definition 2.4 (Def.2.9 of [8]). For an order n twisted Whitney tower \mathcal{W} , let $t_n(\mathcal{W})$ denote the sub-multiset $t_n(\mathcal{W}) \subset t(\mathcal{W})$ consisting of all order n framed trees and order $n/2$ twisted trees in $t(\mathcal{W})$. Define the order n twisted intersection invariant:

$$\tau_n^\infty(\mathcal{W}) := [t_n(\mathcal{W})] \in \mathcal{T}_n^\infty$$

If \mathcal{W} is an order n twisted Whitney tower, then the intersection forest $t(\mathcal{W})$ may a priori contain framed trees of order $> n$ and ∞ -trees of order $> n/2$, but in fact any such Whitney disks in \mathcal{W} can be deleted and/or modified yielding $t_n(\mathcal{W}) = t(\mathcal{W})$ (see Exercise 2.13.1).

Theorem 2.5 (Thm.1.9 of [8]). *A link $L \subset S^3$ bounds an order n twisted $\mathcal{W} \subset B^4$ with $\tau_n^\infty(\mathcal{W}) = 0 \in \mathcal{T}_n^\infty$ if and only if L bounds an order $n + 1$ twisted Whitney tower.*

Idea of proof: For the “only if” direction, first realize relations by geometric constructions, as discussed in section 2.2, to arrange that all trees in $t_n(\mathcal{W})$ occur in oppositely-signed isomorphic pairs. Then use controlled maneuvers to arrange that all order n intersections admit Whitney disks and all Whitney disks of order $\leq n/2$ are framed. See section 4.6 for an outline of this proof, and [8, Sec.4] for details. For the “if” direction see Exercise 2.13.2.

2.4. Quick review of first non-vanishing Milnor invariants.

Let $L \subset S^3$ be an m -component link with fundamental group $G = \pi_1(S^3 \setminus L)$. By [23, Thm.4], if the longitudes of L lie in the $(n + 1)$ -th term G_{n+1} of the lower central series of G , then a choice of meridians induces an isomorphism $\frac{G_{n+1}}{G_{n+2}} \cong \frac{F_{n+1}}{F_{n+2}}$, where $F = F(m)$ is the free group on $\{x_1, x_2, \dots, x_m\}$, with x_i corresponding to the i th meridian.

Let $L = L(m)$ denote the free Lie algebra (over \mathbb{Z}) on generators $\{X_1, X_2, \dots, X_m\}$. It is \mathbb{N} -graded, $L = \bigoplus_n L_n$, where the degree n part L_n is the additive abelian group of length n brackets, modulo Jacobi identities and self-annihilation relations $[X, X] = 0$. The multiplicative abelian group $\frac{F_{n+1}}{F_{n+2}}$ of length $n + 1$ commutators is isomorphic to L_{n+1} , with x_i mapping to X_i and commutators mapping to Lie brackets.

In this setting, denote by l_i the image in L_{n+1} of any i -th longitude of L under the above isomorphisms (l_i is well-defined since choices of longitudes differ only by conjugation), and define the *order n Milnor invariant* $\mu_n(L)$ by

$$\mu_n(L) := \sum_{i=1}^m X_i \otimes l_i \in L_1 \otimes L_{n+1}$$

This definition of $\mu_n(L)$ is the first non-vanishing “total” Milnor invariant of order n , and corresponds to *all* Milnor invariants of length $n + 2$ in the original formulation of [22, 23]. The original $\bar{\mu}$ -invariants are computed from the longitudes via the Magnus expansion as integers modulo indeterminacies coming from invariants of shorter length. Since we will only be concerned with first non-vanishing μ -invariants, we do not use the “bar” notation $\bar{\mu}$.

It turns out that $\mu_n(L)$ lies in the kernel D_n of the bracket map $L_1 \otimes L_{n+1} \rightarrow L_{n+2}$ (e.g. by “cyclic symmetry” [14]).

2.5. The summation maps η_n

The connection between $\tau_n^\infty(\mathcal{W})$ and $\mu_n(L)$ is via a homomorphism $\eta_n : \mathcal{T}_n^\infty \rightarrow D_n$ which is most easily described by regarding rooted trees of order n as elements of L_{n+1} in the usual way: For v a univalent vertex of an order n framed tree t , denote by $B_v(t) \in L_{n+1}$ the Lie bracket of generators X_1, X_2, \dots, X_m determined by the formal bracketing of indices which is gotten by considering v to be a root of t .

Definition 2.6. *Denoting the label of a univalent vertex v by $l(v) \in \{1, 2, \dots, m\}$, the map $\eta_n : \mathcal{T}_n^\infty \rightarrow L_1 \otimes L_{n+1}$ is defined on generators by*

$$\eta_n(t) := \sum_{v \in \mathfrak{t}} X_{l(v)} \otimes B_v(t) \quad \text{and} \quad \eta_n(J^\infty) := \frac{1}{2} \eta_n(\langle J, J \rangle)$$

The first sum is over all univalent vertices v of t , and the second expression lies in $L_1 \otimes L_{n+1}$ because the coefficients of $\eta_n(\langle J, J \rangle)$ are even. Here J is a rooted tree of order j for $n = 2j$. (Recall that J^∞ is obtained from J by labeling the root vertex of J with the ∞ symbol.)

Examples of η_n for $n = 1, 2$:

$$\begin{aligned} \eta_1(1 \leftarrow \frac{3}{2}) &= X_1 \otimes \leftarrow \frac{3}{2} + X_2 \otimes 1 \leftarrow 3 + X_3 \otimes 1 \leftarrow 2 \\ &= X_1 \otimes [X_2, X_3] + X_2 \otimes [X_3, X_1] + X_3 \otimes [X_1, X_2]. \end{aligned}$$

And,

$$\begin{aligned} \eta_2\left(\infty \longleftarrow \begin{array}{c} 2 \\ 1 \end{array}\right) &= \frac{1}{2} \eta_2\left(\begin{array}{c} 1 \\ 2 \end{array} \rightrightarrows \begin{array}{c} 2 \\ 1 \end{array}\right) \\ &= X_1 \otimes \begin{array}{c} 2 \\ 2 \end{array} \rightrightarrows \begin{array}{c} 2 \\ 1 \end{array} + X_2 \otimes \begin{array}{c} 1 \\ 1 \end{array} \rightrightarrows \begin{array}{c} 2 \\ 1 \end{array} \\ &= X_1 \otimes [X_2, [X_1, X_2]] + X_2 \otimes [[X_1, X_2], X_1]. \end{aligned}$$

The image of η_n is equal to the bracket kernel $D_n \subset L_1 \otimes L_{n+1}$, by [11, Lem.32].

Theorem 2.7 ([8]). *If L bounds a twisted Whitney tower \mathcal{W} of order n , then the order q Milnor invariants $\mu_q(L)$ vanish for $q < n$, and*

$$\mu_n(L) = \eta_n \circ \tau_n^\infty(\mathcal{W}) \in D_n$$

Idea of proof: The existence of the order n twisted \mathcal{W} implies, via Dwyer’s Theorem, that the inclusion $S^3 \setminus L \rightarrow B^4 \setminus \mathcal{W}$ induces an isomorphism on the $(n+1)$ th lower central quotients of π_1 , so the longitudes of L can be computed in $B^4 \setminus \mathcal{W}$. It turns out that the corresponding iterated commutators are displayed exactly according to $\eta_n \circ \tau_n^\infty(\mathcal{W})$, with the key observation being that a meridian to a Whitney disk is a commutator of meridians to the sheets paired by the Whitney disk. See [11, Sec.4] or [3, Thm.3.1].

2.6. The order n twisted Whitney tower filtration on links

Recall that a link $L \subset S^3 = \partial B^4$ bounds an order n twisted Whitney tower \mathcal{W} if $\mathcal{W} \subset B^4$ is an order n twisted Whitney tower whose order 0 surfaces are immersed disks bounded by the components of L .

We say that links L_0 and L_1 in S^3 are *twisted Whitney tower concordant of order n* if $L_0 \subset S^3 \times \{0\}$ and $L_1 \subset S^3 \times \{1\}$ cobound a collection $A \looparrowright S^3 \times [0, 1]$ of immersed annuli such that A supports an order n twisted Whitney tower (with A inducing the reversed orientation on L_1).

Denote by W_n^∞ the set of links in S^3 bounding order n twisted Whitney towers in B^4 modulo the equivalence relation of order $n+1$ twisted Whitney tower concordance.

The twisted “order-raising” Theorem 2.5 implies the following essential criterion for links to represent equal elements in W_n^∞ :

Corollary 2.8 ([8, Cor.3.3]). *Links L_0 and L_1 represent the same element of W_n^∞ if and only if there exist order n twisted Whitney towers \mathcal{W}_i in B^4 with $\partial \mathcal{W}_i = L_i$ and $\tau_n^\infty(\mathcal{W}_0) = \tau_n^\infty(\mathcal{W}_1) \in \mathcal{T}_n^\infty$.*

Proof. If L_0 and L_1 are equal in W_n^∞ then they cobound A supporting an order $n+1$ twisted Whitney tower \mathcal{V} in $S^3 \times I$, and any order n twisted Whitney tower \mathcal{W}_1 in B^4 bounded by L_1 can be extended by \mathcal{V} to form an order n twisted Whitney tower \mathcal{W}_0 in B^4 bounded by L_0 , with $\tau_n^\infty(\mathcal{W}_0) = \tau_n^\infty(\mathcal{W}_1) \in \mathcal{T}_n^\infty$ since $\tau_n^\infty(\mathcal{V})$ vanishes.

Conversely, suppose that L_0 and L_1 bound order n twisted Whitney towers \mathcal{W}_0 and \mathcal{W}_1 in 4–balls B_0^4 and B_1^4 , with $\tau_n^\infty(\mathcal{W}_0) = \tau_n^\infty(\mathcal{W}_1)$. Then constructing $S^3 \times I$ as the connected sum $B_0^4 \# B_1^4$ (along balls in the complements of \mathcal{W}_0 and \mathcal{W}_1), and tubing together the corresponding order zero disks of \mathcal{W}_0 and \mathcal{W}_1 , and taking the union of the Whitney disks in \mathcal{W}_0 and \mathcal{W}_1 , yields a collection A of properly immersed annuli connecting L_0 and L_1 and supporting an order n twisted Whitney tower \mathcal{V} . Since the orientation of the ambient 4–manifold has been reversed for one of the original Whitney towers, say \mathcal{W}_1 , which results in a global sign change for $\tau_n^\infty(\mathcal{W}_1)$, it follows that \mathcal{V} has vanishing order n intersection invariant:

$$\tau_n^\infty(\mathcal{V}) = \tau_n^\infty(\mathcal{W}_0) - \tau_n^\infty(\mathcal{W}_1) = \tau_n^\infty(\mathcal{W}_0) - \tau_n^\infty(\mathcal{W}_0) = 0 \in \mathcal{T}_n^\infty$$

So by Theorem 2.5, A is homotopic (rel ∂) to A' supporting an order $n+1$ twisted Whitney tower, and hence L_0 and L_1 are equal in W_n^∞ . \square

The *band sum* $L \#_\beta L' \subset S^3$ of oriented m -component links L and L' along bands β is defined as follows: Form S^3 as the connected sum of 3–spheres containing L and L' along balls in the link complements. Let β be a collection of disjointly embedded oriented bands joining

like-indexed link components such that the band orientations are compatible with the link orientations. Take the usual connected sum of each pair of components along the corresponding band. Although it is well-known that the concordance class of $L\#_{\beta}L'$ depends in general on β , it turns out that the image of $L\#_{\beta}L'$ in W_n^{∞} does not depend on β :

Lemma 2.9 ([8]). *For links L and L' representing elements of W_n^{∞} , any band sum $L\#_{\beta}L'$ represents an element of W_n^{∞} which only depends on the equivalence classes of L and L' in W_n^{∞} .*

See Lemma 3.6 of [8] or Exercise 2.13.4 for a proof of Lemma 2.9 using the following:

Lemma 2.10 ([8]). *If L and L' bound order n twisted Whitney towers \mathcal{W} and \mathcal{W}' in B^4 , then for any bands β there exists an order n twisted Whitney tower $\mathcal{W}^{\#} \subset B^4$ bounded by $L\#_{\beta}L'$, such that $t(\mathcal{W}^{\#}) = t(\mathcal{W}) + t(\mathcal{W}')$.*

See Lemma 3.7 of [8], or Exercise 2.13.3, for a proof.

2.7. Definition of the realization maps

We define “realization” maps $R_n^{\infty} : \mathcal{T}_n^{\infty} \rightarrow W_n^{\infty}$ for all n as follows: Given any group element $g \in \mathcal{T}_n^{\infty}$, by Lemma 2.11 just below there exists an m -component link $L \subset S^3$ bounding an order n twisted Whitney tower $\mathcal{W} \subset B^4$ such that $\tau_n^{\infty}(\mathcal{W}) = g \in \mathcal{T}_n^{\infty}$. Define $R_n^{\infty}(g)$ to be the class determined by L in W_n^{∞} . This is well-defined (does not depend on the choice of such L) by Corollary 2.8.

Lemma 2.11 ([8]). *For any multiset $\sum_p \epsilon_p \cdot t_p + \sum_j \omega(W_j) \cdot J^{\infty}$ there exists a link L bounding a Whitney tower \mathcal{W} with intersection forest $t(\mathcal{W}) = \sum_p \epsilon_p \cdot t_p + \sum_j \omega(W_j) \cdot J^{\infty}$.*

The proof of Lemma 2.11 follows Tim Cochran’s technique of “Bing-doubling along a tree” to realize individual trees, and then uses band sums of links via Lemma 2.10 to realize sums (multiset unions) of trees (see Lemma 3.8 of [8] or Exercise 2.13.5).

For instance, to compute the image of the framed tree $t = \langle (1, 2), (3, 1) \rangle = \frac{1}{2} \triangleright \frac{1}{3} \leftarrow$ under R_2^{∞} , one Bing-doubles each component of the Hopf link and then bands together two components as in Figure 2.4 to get a link L bounding a Whitney tower \mathcal{W} with $\tau_2^{\infty}(\mathcal{W}) = t$. Any framed tree can be realized analogously by applying iterated Bing-doubling to the Hopf link to get the desired tree shape, and then banding components to get the desired univalent labels.




Figure 2.4

To compute the image of the twisted tree $(2, 1)^{\infty} = \frac{2}{1} \triangleright \infty$ under R_2^{∞} , one applies a twisted Bing-double to the unknot as in Figure 2.5 to get L bounding a Whitney tower \mathcal{W} with $\tau_2^{\infty}(\mathcal{W}) = (2, 1)^{\infty}$. Any twisted tree J^{∞} with twisting coefficient $\omega \in \mathbb{Z}$ can be realized similarly by starting with a single ω -twisted Bing-doubling of the unknot and then applying iterated

(untwisted) Bing-doublings to get the desired tree shape, and then banding components to get the desired univalent labels.




Figure 2.5

In combination with Theorem 2.7, the following result will play a key role in classifying order n twisted Whitney towers in the 4-ball:

Theorem 2.12 ([8]). *The realization maps are epimorphisms $R_n^\infty : \mathcal{T}_n^\infty \rightarrow W_n^\infty$.*

Proof. From Lemma 2.9 the band sum of links gives a well-defined operation in W_n^∞ which is clearly associative and commutative, with the m -component unlink representing an identity element. The realization maps are homomorphisms by Lemma 2.10 and surjectivity is proven as follows: Given any link L representing an element of W_n^∞ , choose a twisted Whitney tower \mathcal{W} of order n with boundary L and compute $\tau := \tau_n^\infty(\mathcal{W})$. Then take $L' := R_n^\infty(\tau)$, a link that's in the image of R_n^∞ and for which we know a Whitney tower \mathcal{W}' with boundary L' and $\tau^\infty(\mathcal{W}') = \tau$. By Corollary 2.8 it follows that L and L' represent the same element in W_n^∞ . \square

2.8. Computing the order n twisted Whitney tower filtration

From Theorem 2.7 and Theorem 2.12 we have the following commutative triangle diagram ∇ of epimorphisms:

$$(\nabla) \quad \begin{array}{ccc} \mathcal{T}_n^\infty & \xrightarrow{R_n^\infty} & W_n^\infty \\ & \searrow \eta_n & \downarrow \mu_n \\ & & D_n \end{array}$$

The following algebraic result is a consequence of the proof [9] of a combinatorial conjecture of J. Levine:

Theorem 2.13 ([8]). *The maps $\eta_n : \mathcal{T}_n^\infty \rightarrow D_n$ are isomorphisms for $n \equiv 0, 1, 3 \pmod{4}$.*

From this theorem we immediately get:

Corollary 2.14. *For $n \equiv 0, 1, 3 \pmod{4}$, the Milnor invariants $\mu_n : W_n^\infty \rightarrow D_n$ and the twisted realization maps $R_n^\infty : \mathcal{T}_n^\infty \rightarrow W_n^\infty$ are isomorphisms.*

By [24], D_n is a free abelian group of known rank for all n , so we have a complete computation of $W_n^\infty \cong D_n \cong \mathcal{T}_n^\infty$ in three quarters of the cases.


Remark 2.15. *It also follows that in these orders the order n twisted intersection invariant $\tau_n^\infty(\mathcal{W}) \in \mathcal{T}_n^\infty$ only depends on the concordance class of $L = \partial\mathcal{W}$. In particular, $\tau_n^\infty(\mathcal{W}) \in \mathcal{T}_n^\infty$ is independent of the choice of Whitney tower \mathcal{W} on the (unique) homotopy class rel ∂ of the order 0 immersed disks bounded by L .*

Towards understanding the remaining cases $n \equiv 2 \pmod{4}$, we have another consequence of [9] which was derived as Corollary 6.6 of [8]:

Proposition 2.16 ([8]). For rooted trees J of order $j - 1$, the map $1 \otimes J \mapsto \infty \leftarrow J \in \mathcal{T}_{4j-2}^\infty$ induces an isomorphism:

$$\mathbb{Z}_2 \otimes L_j \cong \text{Ker}\{\eta_{4j-2} : \mathcal{T}_{4j-2}^\infty \rightarrow D_{4j-2}\}$$

So $\mathbb{Z}_2 \otimes L_j$ is also an upper bound on $K_{4j-2}^\infty := \text{Ker}\{\mu_{4j-2} : W_{4j-2}^\infty \rightarrow D_{4j-2}\}$. Denoting by $\alpha_j^\infty : \mathbb{Z}_2 \otimes L_j \twoheadrightarrow K_{4j-2}^\infty$ the epimorphism induced by R_{4j-2}^∞ , we have the following extension of the above triangle diagram ∇ in these orders:



By inverting the induced isomorphism $\overline{\alpha_j^\infty}$ on $(\mathbb{Z}_2 \otimes L_j) / \text{Ker } \alpha_j^\infty$, we get the following definition:

Definition 2.17. The higher-order Arf invariants Arf_j are defined by

$$\text{Arf}_j := (\overline{\alpha_j^\infty})^{-1} : K_{4j-2}^\infty \rightarrow (\mathbb{Z}_2 \otimes L_j) / \text{Ker } \alpha_j^\infty$$

As a corollary we get the computation of W_n^∞ for all n , and a characterization of links bounding order n twisted Whitney towers:

Corollary 2.18 ([8]). The abelian groups W_n^∞ are classified by Milnor invariants μ_n and, in addition, higher-order Arf invariants Arf_j for $n = 4j - 2$.

In particular, a link bounds an order $n + 1$ twisted \mathcal{W} if and only if its Milnor invariants and higher-order Arf invariants vanish up to order n .


2.9. The higher-order Arf invariant Conjecture

In the case $j = 1$ of Definition 2.17, $\text{Ker } \alpha_1^\infty$ is trivial and we have $\text{Arf}_1 : K_2^\infty \xrightarrow{\cong} (\mathbb{Z}_2 \otimes L_1) \cong (\mathbb{Z}_2)^m$, given by classical Arf invariants of the link components [11, Lem.10].

The higher-order Arf invariant conjecture states that $\text{Ker } \alpha_j^\infty$ is trivial for all j :

Conjecture 2.19. $\text{Arf}_j : K_{4j-2}^\infty \rightarrow \mathbb{Z}_2 \otimes L_j$ are isomorphisms for all j .

Assuming this conjecture the classification in these orders would be described by the following diagram:



Conjecture 2.19 would imply $W_n^\infty \xrightarrow{\tau_n^\infty} \mathcal{T}_n^\infty$ is an isomorphism for all n . That is, the intersection invariants τ_n^∞ taking values in \mathcal{T}_n^∞ are independent of Whitney tower choices and characterize links bounding order n twisted Whitney towers.




Figure 2.6. The (untwisted) Bing double of the Figure-8 knot ($L_1 \cup L_2 \subset S^3$ on left) bounds an order 6 twisted Whitney tower $D_1 \cup D_2 \cup W_{(1,2)} \cup W_{((1,2),(1,2))} \subset B^4$, with $\omega(W_{((1,2),(1,2))}) = 1$. Here the order zero disks D_i bounded by the L_i are contained near S^3 in a collar (not shown) which corresponds to the left arrow, and the intersection pair $D_1 \cap D_2$ which is paired by $W_{(1,2)}$ corresponds to two “unclasp” crossing changes in a null-homotopy of $L_1 \cup L_2$ supported in this collar. Thus, $\partial W_{(1,2)}$ is a parallel of the Figure-8 knot core of the band used to create the Bing double.

2.10. Problems, Questions and re-formulations

As just mentioned in section 2.9, Conjecture 2.19 is true for $k = 1$, with Arf_1 given by the classical Arf invariants of the link components [11], but it remains an open problem whether Arf_k is non-trivial for any $k > 1$. The links $R_{4k-2}^\infty(\infty \leftarrow J)$ realizing the image of Arf_k can all be constructed as internal band sums of iterated Bing doubles of knots having non-trivial classical Arf invariant ([11, Lem.13]), see Figure 2.6. Such links are known not to be slice by work of J.C. Cha [2], providing evidence in support of Conjecture 2.19.

So the fundamental open problem in this setting is:

Open Problem 2.20. Determine the precise image of $\text{Arf}_j \leq \mathbb{Z}_2 \otimes L_j$ for $j \geq 2$.

The following specific lowest order open question is already interesting:

Question 2.21. Does the Bing double of the Figure-8 knot $R_6^\infty(\infty \leftarrow \begin{smallmatrix} (1,2) \\ (1,2) \end{smallmatrix}) \in W_6^\infty$, shown in Figure 2.6, bound an order 7 twisted Whitney tower?

This importance of this lowest order open question is magnified by the following fact [11, Prop.14]:

Proposition 2.22 ([11]). If the Bing double of the Figure-8 knot does bound an order 7 twisted Whitney tower, then Arf_j are trivial for all $j \geq 2$.

Conjecture 2.19 predicts a negative answer to Question 2.21, which can be phrased as the following restriction on the possible twisted Whitney towers on 2-spheres in the 4-ball:

Conjecture 2.23. There does not exist $A : S^2 \cup S^2 \rightarrow B^4$ supporting \mathcal{W} with

$$t(\mathcal{W}) = \infty \leftarrow \begin{smallmatrix} (1,2) \\ (1,2) \end{smallmatrix} \quad (\text{possibly } + \text{ higher-order trees}).$$

2.11. Higher-order Arf invariant Conjecture and Finite Type invariants

Habegger and Masbaum [16] have shown that Milnor invariants are the only rational finite type concordance invariants of (string) links. The classical Arf invariant of a knot is known to

be determined by the first non-trivial finite type isotopy invariant, and as stated above, Arf_1 corresponds to the classical Arf invariants of the link components. So it is natural to ask:

Question 2.24. *Are the Arf_j for $j > 1$ also determined by finite type isotopy invariants?*

For the first interesting case in this setting, a negative answer to Question 2.21 can be formulated as:

Conjecture 2.25. *The sum of trees in Figure 2.7 represents a non-trivial finite type concordance invariant of 2-component string links (first-non-vanishing, $\mathbb{Z}/2\mathbb{Z}$ -coefficients).*




Figure 2.7

The invariant of Conjecture 2.25 would be finite type *degree* 6, and all finite type concordance invariants of string links in degrees ≤ 5 have been characterized by J-B. Meilhan and A. Yasuhara in [21]. So this conjecture appears to lie at the frontier of current understanding in this setting.

Degree 6 finite type invariants are related to order 5 intersection invariants in the setting of *framed* Whitney towers [7, 30], and the sum of trees in Figure 2.7 (modulo higher-order trees) can be gotten by applying twisted IHX moves and boundary twists to a Whitney tower \mathcal{W} with $t(\mathcal{W}) = \infty \leftarrow \begin{smallmatrix} (1,2) \\ (1,2) \end{smallmatrix}$ bounded by any link $R_6^\infty(\infty \leftarrow \begin{smallmatrix} (1,2) \\ (1,2) \end{smallmatrix}) \in W_6^\infty$, see Exercise 2.13.12. This illustrates how the higher-order Arf invariants shift down one order in the *framed* order n Whitney tower filtration [8]. Conjecture 2.25 corresponds in this setting to:

Conjecture 2.26. *The Bing double of any knot with non-trivial classical Arf invariant does not bound an order 6 framed Whitney tower.*

By Exercise 2.13.15, Conjecture 2.26 can also be phrased as a restriction on Whitney towers supported by 2-spheres in 4-space:

Conjecture 2.27. *There does not exist $A : S^2 \cup S^2 \looparrowright B^4$ supporting \mathcal{W} with $t(\mathcal{W})$ equal to the trees in Figure 2.7, possibly plus higher-order trees.*

We remark that each of the two trees in Figure 2.7 *individually* represents a non-trivial *higher-order Sato–Levine invariant* which is determined by a non-trivial order 6 μ -invariant [8], so neither of these trees can appear by itself (plus higher-order trees) as $t(\mathcal{W})$ for \mathcal{W} on 2-spheres in B^4 (compare Exercise 2.13.10).

Recent work of Danica Kosanović [19] includes progress towards showing that the *embedding calculus* of Goodwillie–Weiss [15, 34] determines universal finite type invariants for knots over the integers, as conjectured in [1]. Kosanović works with certain 2-complexes called *capped gropes* which are very closely related to Whitney towers [8, 26], and one could hope that, via a generalization of [19] to links, the homotopy-theoretic techniques of the embedding calculus might be able to detect the higher-order Arf invariants:

Question 2.28. *Can the Arf_j be detected by the embedding calculus?*

2.12. Higher-order Arf invariant conjecture and transfinite Milnor invariants

In [4] Jae Choon Cha and Kent Orr defined certain transfinite invariants of 3-manifolds which can be interpreted as providing a generalization of Milnor invariants. Their invariants include

certain finite-length θ_k -invariants which can take values in finite abelian groups, and in [4, Sec.14(9)] they ask:

Question 2.29. *Are the higher-order Arf invariants related to the θ_k -invariants?*

2.13. Section 2 Exercises

2.13.1. Exercise

If \mathcal{W} is an order n twisted Whitney tower, then the intersection forest $t(\mathcal{W})$ may contain framed trees of order $> n$ and ∞ -trees of order $> n/2$ in addition to those representing $\tau_n^\infty(\mathcal{W}) := [t_n(\mathcal{W})] \in \mathcal{T}_n^\infty$. Show that by deleting Whitney disks of order $> n$, boundary-twisting, and pushing-down intersections (Figure 1.26), these higher-order trees in $t(\mathcal{W})$ can be eliminated without changing $\tau_n^\infty(\mathcal{W})$ while preserving the twisted order n of the resulting Whitney tower. (See discussion in [8, Sec.4.1, 'Notation and Conventions'.])

2.13.2. Exercise

If \mathcal{W} is an order $n + 1$ twisted Whitney tower, observe that by definition \mathcal{W} is also an order n twisted Whitney tower, and check that $\tau_n^\infty(\mathcal{W}) = 0 \in \mathcal{T}_n^\infty$.

2.13.3. Exercise

Show that if L and L' bound order n twisted Whitney towers \mathcal{W} and \mathcal{W}' in B^4 , then for any β there exists an order n twisted Whitney tower $\mathcal{W}^\# \subset B^4$ bounded by $L \#_\beta L'$, such that $t(\mathcal{W}^\#) = t(\mathcal{W}) + t(\mathcal{W}')$. (Lemma 3.7 of [8].)

2.13.4. Exercise

Use the previous exercise to show that for links L and L' representing elements of W_n^∞ , any band sum $L \#_\beta L'$ represents an element of W_n^∞ which only depends on the equivalence classes of L and L' in W_n^∞ . (Lemma 3.6 of [8].)

2.13.5. Exercise

Given any framed tree $\langle I, J \rangle$, construct a link $L \subset S^3$ bounding $\mathcal{W} \subset B^4$ with $t(\mathcal{W}) = \langle I, J \rangle$. HINT: Apply Bing-doubling as needed to the Hopf link, as in Figure 2.4. (See Lemma 3.8 of [8].)

2.13.6. Exercise

Given any integer n and any rooted tree J of positive order, construct a link $L \subset S^3$ bounding $\mathcal{W} \subset B^4$ with $t(\mathcal{W}) = n \cdot J^\infty$. HINT: Start with the n -twisted Bing double of the unknot (see Figure 2.5 for the case $n = 1$), then apply untwisted iterated Bing-doubling as needed. (See Lemma 3.8 of [8].)

2.13.7. Exercise

From Exercise 1.22.9, the Bing-double of the Hopf link bounds an order 2 twisted Whitney tower \mathcal{W} as in Figure 1.9. Show that the Bing-double of the Hopf link does not bound an order 3 twisted Whitney tower by checking that the order 2 Milnor invariant $\mu_2 = \eta_2 \circ \tau_2^\infty(\mathcal{W}) \in L_1 \otimes L_3$ is non-zero. Conclude that the Bing-double of the Hopf link also does not bound an order 3 framed Whitney tower (since for any n an order n framed Whitney tower is also an order n twisted Whitney tower by definition).

2.13.8. Exercise

Compute the order 2 Milnor invariants $\mu_2 = \eta_2 \circ \tau_2^\infty(\mathcal{W})$ for the links in Figures 2.4 and 2.5.

2.13.9. Exercise

Order 0 twisted tree groups and Milnor invariants:

Check from the definitions in section 2.4 that for $i \neq j$ the coefficient of $X_i \otimes X_j$ in $\mu_0(L)$ is the linking number of L_i and L_j , which, via the well-known computation of linking numbers as the count of signed intersections between the properly immersed disks D_i and D_j bounded by L_i and L_j , is also equal to the coefficient of $X_i \otimes X_j$ in $\eta_0(\tau_0^\infty(\mathcal{W}))$.

Although Milnor invariants are not usually defined for knots, for framed links it is natural to consider the framing f_i of L_i as an order 0 (length 2) integer Milnor invariant, and the coefficient of $X_i \otimes X_i$ in $\mu_0(L)$ is exactly f_i when this framing is used to determine the i th longitude. To see that the coefficient in $\eta_0(\tau_0^\infty(\mathcal{W}))$ of $X_i \otimes X_i$ is also equal to f_i , let d_i denote the number of positive self intersections of D_i minus the number of negative self-intersections of D_i . Then the relative Euler number of D_i with respect to the framing f_i on $L_i = \partial D_i$ is equal to $f_i - 2d_i$ (see e.g. Figure 19 of [8] and accompanying discussion), and the terms of $\tau_0^\infty(\mathcal{W})$ which contribute via η_0 to the coefficient of $X_i \otimes X_i$ are exactly $(d_i) \cdot i \longleftarrow i + (f_i - 2d_i) \cdot \infty \longrightarrow i$, which get sent by η_0 to $(f_i) \cdot X_i \otimes X_i$.

2.13.10. Exercise

Show that if $g \in \mathcal{T}_n^\infty$ is such that $\eta_n(g) \neq 0 \in D_n$, then there does not exist any twisted Whitney tower \mathcal{W} on 2-spheres in B^4 such that $t(\mathcal{W})$ represents $g \in \mathcal{T}_n^\infty$. (HINT: Otherwise tubing the spheres into disks bounded by an unlink would “create” non-trivial μ -invariants.)

2.13.11. Exercise

Check that $\eta_6(\text{tree}((1, 2), 1), 2)^\infty \neq 0 \in L_1 \otimes L_6$. Conclude from Exercise 2.13.10 that there does not exist $A : S^2 \cup S^2 \looparrowright B^4$ supporting a Whitney tower \mathcal{W} with $t(\mathcal{W}) = \text{tree}((1, 2), 1), 2)^\infty$, perhaps plus trees of order > 6 .

2.13.12. Exercise

Use the Whitney move twisted IHX relation of Lemma 4.5 and boundary-twisting to get the two trees in Figure 2.7 (plus higher-order trees) from the single tree $\text{tree}((1, 2), (1, 2))^\infty$.

2.13.13. Exercise

Check that $\eta_2 \circ \tau_2^\infty(\mathcal{W})$ vanishes for the Bing double of Figure-8 knot.

2.13.14. Exercise

If a knot K bounds a Whitney tower \mathcal{W} with $t(\mathcal{W}) = (1, 1)^\infty$, show that the Bing-double of K bounds a Whitney tower \mathcal{V} with $t(\mathcal{V}) = ((1, 2), (1, 2))^\infty$.

2.13.15. Exercise

Check that Conjecture 2.26 and Conjecture 2.27 are equivalent.

2.13.16. Exercise

Show that if there exists a pair of 2-spheres in S^4 supporting a Whitney tower \mathcal{W} such that $t(\mathcal{W}) = ((1, 2), (1, 2))^\infty$, then all higher order Arf invariants Arf_j for $j \geq 2$ vanish on all links. (See [11, Prop.14].)

3. Whitney towers on 2-spheres in 4-manifolds

In this section we consider Whitney towers on $A = A_1, A_2, \dots, A_m \looparrowright X$, where each A_i is a 2-sphere, and X is a 4-manifold. All manifolds are oriented and based. Each A_i is equipped with a *whisker*, i.e. an arc running between the basepoint of A_i and the basepoint of X .

We will assume that A is generically immersed, and each sphere A_i has the same number of positive self-intersections as negative self-intersections. This can always be arranged by performing *cuspidal homotopies* ([13, Chap.1], same as the interior twist operation of Figure 2.3), and is a natural assumption in the setting of Whitney towers since it is satisfied as soon as A supports a (framed or twisted) Whitney tower of positive order. Regular homotopy classes of such immersions are in one-to-one correspondence with homotopy classes, see e.g. [25, Thm.1.2]. Up to isotopy, a regular homotopy is a sequence of finger moves and Whitney moves, so homotopy invariance of invariants defined from the intersection forests of Whitney towers can be checked combinatorially (e.g. section 4.1.9).

The classical intersection invariant $\lambda(A_i, A_j)$ and self-intersection invariant $\mu(A_i)$ are recalled in section 3.1, and we formulate these invariants in the language of Whitney towers as order 0 invariants which give the complete obstruction to A supporting an order 1 framed Whitney tower in section 3.2. After explaining how the order 0 intersection pairing gives the complete obstruction to “pulling apart” a pair of spheres, i.e. making them disjoint by a homotopy (section 3.3), edge decorations in $\pi_1 X$ for order 1 trees are introduced in section 3.4 and the order 1 non-repeating invariant generalizing the classical intersection pairing is defined in the setting of pulling apart triples of components of A (section 3.5 and Theorem 3.2). In section 3.6 we briefly touch on possible higher-order non-repeating invariants and a general obstruction theory for pulling apart multiple components.

The order 1 generalization of the classical self-intersection invariant is defined in section 3.7. Here twistings on Whitney disks are relevant, and the vanishing of the invariant is equivalent to the existence of an order 2 framed Whitney tower on A (Theorem 3.4).

Throughout this section, relevant open questions, problems and conjectures are included in the discussions. Proofs of the main Theorems 3.2 and 3.4 are given in Section 4.

3.1. Classical intersection form

The first obstructions to making the components of A pairwise disjointly embedded by a homotopy are the *intersection invariants* $\lambda(A_i, A_j)$, which take values in the integral fundamental group ring $\mathbb{Z}[\pi_1 X]$, and the *self-intersection invariants* $\mu(A_i)$, which take values in a quotient of $\mathbb{Z}[\pi_1 X]$. These invariants are defined as follows. Associate to each transverse intersection point $p \in A_i \pitchfork A_j$ the element $g_p \in \pi_1 X$ determined by a loop through A_i and A_j which changes sheets at p . More precisely, such a *sheet-changing loop* runs along the whisker for A_i , then along any choice of path in A_i to p , then along any choice of path in A_j to the whisker of A_j , then along the whisker for A_j . Sheet-changing loops are required to avoid all singularities of A_i and A_j other than p , so g_p does not depend on the choices of paths in A_i and A_j between p and their whiskers because the domains of A_i and A_j are simply connected. See Figure 3.1.

Summing over all such intersection points, with the usual notion of the sign $\epsilon_p \in \{+, -\}$, defines:

$$\lambda(A_i, A_j) := \sum_{p \in A_i \pitchfork A_j} \epsilon_p \cdot g_p \in \mathbb{Z}[\pi_1 X]$$



Figure 3.1. $g_p \in \pi_1 X$ for $p \in A_i \pitchfork A_j$ with $i \neq j$ (left), and for $i = j$ (right).

and

$$\mu(A_i) := \sum_{p \in A_i \pitchfork A_i} \epsilon_p \cdot g_p \in \frac{\mathbb{Z}[\pi_1 X]}{\{g - g^{-1}\}}.$$

The relations in $\mathbb{Z}[\pi_1 X]$ induced by $g = g^{-1} \in \pi_1 X$ account for choices of orientations on the sheet-changing loops through the self-intersections of A_i .

Up to isotopy, a regular homotopy of A is a sequence of finger moves and Whitney moves, each of which either introduces or eliminates oppositely-signed intersections having the same group element. So λ and μ are invariant under regular homotopy, e.g. [25, Sec.4].

3.2. Order 0 invariants

Here we slightly reformulate the classical invariants λ and μ in the language of Whitney towers, with an eye towards higher-order generalizations.

The union $A = A_1 \cup A_2 \cup \dots \cup A_m \subset X$ is by Definition 1.7 a framed Whitney tower of order 0. To each order 0 intersection $p \in A_i \pitchfork A_j$ is associated the order 0 tree $t_p = \langle i, j \rangle$, as in section 1.10, and we think of $t_p \subset A_i \cup A_j$ as an embedded sheet-changing edge near p with one univalent vertex in A_i and the other univalent vertex in A_j . For each such t_p , choose a path in A_i from the basepoint of A_i to the i -labeled univalent vertex of t_p , and a path in A_j from the j -labeled univalent vertex of t_p to the basepoint of A_j . The union of t_p (oriented from i to j) together with these paths and the whiskers on A_i and A_j , defines a sheet-changing loop representing $g_p \in \pi_1 X$, just as in section 3.1. We call the tree t_p , together with the label g_p on its edge and an orientation of the edge from i to j , a *decorated tree* for p .

Let \mathcal{T}_0 denote the quotient of the free abelian group on order 0 decorated trees by the following OR *orientation* relation:

$$\text{OR: } i \xrightarrow{g_p} j = i \xleftarrow{g_p^{-1}} j$$

Now the classical invariants $\lambda_0 := \lambda$ and $\mu_0 := \mu$ can be expressed as a single order 0 invariant $\tau_0(A)$ represented by the intersection forest $t(A)$:

$$\tau_0(A) := \sum_{p \in A_i \pitchfork A_j} \epsilon_p \cdot i \xrightarrow{g_p} j \in \mathcal{T}_0$$

In the language of Whitney towers we have:

Theorem 3.1. $\tau_0(A) = 0$ if and only if A supports an order 1 framed Whitney tower.

Recall from Definition 1.7 that A supporting an order 1 framed Whitney tower \mathcal{W} means that all intersections of A are paired by framed order 1 Whitney disks in \mathcal{W} . The idea of the “only if” direction of the proof is that all intersections can be arranged in oppositely-signed pairs having the same group element, after orienting sheet-changing loops appropriately.

Then null-homotopic Whitney circles can be constructed from the pairs of sheet-changing loops. These Whitney circles bound immersed Whitney disks which can be made framed by boundary-twisting (Figure 2.2). See Exercise 3.9.5. For the “if” direction see Exercise 3.9.4. See also [25, Lem.4.3].

It can always be arranged that the Whitney disks in an order 1 framed Whitney tower are disjointly embedded (Exercise 1.22.13), but they will in general have interior intersections with A which obstruct using them to guide Whitney moves homotoping A to an embedding.

We remark that for half-dimensional spheres $A : S_1^d, S_2^d, \dots, S_m^d \looparrowright X^{2d}$ in a $2d$ -dimensional manifold with $d > 2$ the invariants analogous to λ and μ give the complete obstruction to embedding A in X^{2d} , because the interiors of Whitney disks on A will have interiors disjoint from A by general position.

Notice that $\tau_0(A)$ splits into a direct sum of *non-repeating* and *repeating* invariants $\tau_0(A) = \sum_{i \neq j} \lambda_0(A_i, A_j) \oplus \sum_i \tau_0(A_i)$. Before generalizing the full order 0 invariant $\tau_0(A)$ to an order 1 invariant $\tau_1(A)$, we will first consider the intermediate problem of generalizing the non-repeating summands $\lambda_0(A) := \sum_{i \neq j} \lambda_0(A_i, A_j)$ to order 1 invariants $\lambda_1(A_i, A_j, A_k)$ for triples with distinct i, j, k , and discuss the relationship to “pulling apart” triples of components (making them pairwise disjoint by a homotopy).

3.3. Pulling apart pairs of spheres

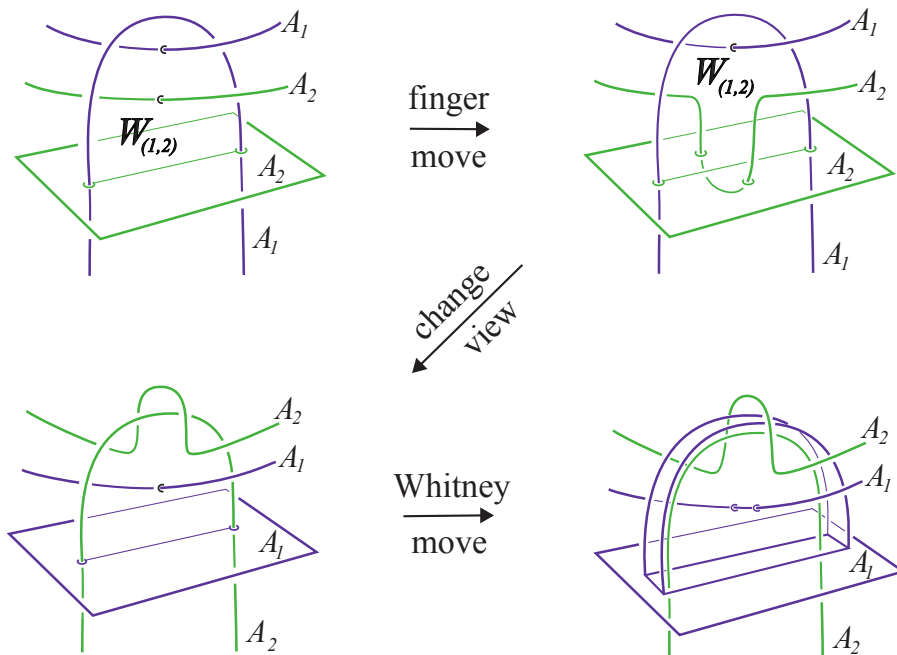


Figure 3.2

The vanishing of $\lambda_0(A_1, A_2) = 0 \in \mathbb{Z}[\pi_1 X]$ implies the existence of Whitney disks pairing $A_1 \pitchfork A_2$ (Exercise 3.9.1), and the union of $A_1 \cup A_2$ together with a collection of such Whitney disks forms an *order 1 non-repeating Whitney tower* \mathcal{W} , cf. Definition 1.9. As an illustration of the proof of Theorem 1.10 in this easiest case, Figure 3.2 shows how the existence of such a \mathcal{W} leads to a homotopy that makes A_1 and A_2 disjoint: Finger moves as in the top of Figure 3.2 make A_2 disjoint from the interiors of all the Whitney disks, at the cost of only creating self-intersections in A_2 . Now doing all the Whitney moves on A_1 makes $A_1 \cap A_2 = \emptyset$ at the cost of only creating self-intersections in A_1 as in the bottom of Figure 3.2. Any self-intersections and intersections among the Whitney disks will only lead to the creation of more self-intersections

in A_1 upon doing the Whitney moves, hence such intersections have been suppressed from view in Figure 3.2.

The procedure of Figure 3.2 appears to fail in the presence of a third sphere A_3 , as it is not clear how to eliminate any $W_{(1,2)} \cap A_3$ without creating more intersections between A_3 and A_1 or A_2 . In order to generalize $\lambda_0(A_1, A_2)$ to an order 1 invariant $\lambda_1(A_1, A_2, A_3)$ which “counts” such order 1 intersections $W_{(i,j)} \cap A_k$ and gives the complete obstruction to pulling apart triples of 2-spheres, we next introduce edge decorations for order 1 trees.

3.4. Decorated trees for order 1 intersections

Let $W_{(i,j)}$ be a Whitney disk for a pair of intersections in $A_i \pitchfork A_j$. To each intersection $p \in W_{(i,j)} \pitchfork A_k$ we associate a *decorated* order 1 tree t_p which is gotten by labeling each edge of the usual tree $\langle (i, j), k \rangle$ from section 1.10 by an element of $\pi_1 X$ as in Figure 3.3. This requires a choice of whisker running from the trivalent vertex of each tree to the basepoint of X , and the group elements are determined by loops formed using sheet-changing edges oriented towards the Whitney disk followed by the chosen whisker emanating from the trivalent vertex. More precisely, taking t_p to be embedded in $A_i \cup A_j \cup A_k \cup W_{(i,j)}$ as in Figure 3.3, each edge is a sheet-changing embedded arc which can be oriented to change sheets into the Whitney disk, and each univalent vertex can be connected by a path in its order 0 surface to the surface basepoint. Together with the whiskers on A_i, A_j, A_k and the trivalent vertex we get three oriented loops determining the corresponding edge decorations $a, b, c \in \pi_1 X$.




Figure 3.3. Edge decorations $a, b, c \in \pi_1 X$.

3.5. The order 1 non-repeating intersection invariant λ_1

Consider a triple of immersed spheres $A = A_1, A_2, A_3 \pitchfork X$ supporting an *order 1 non-repeating Whitney tower* \mathcal{W} , i.e. all intersections $A_i \pitchfork A_j$ for distinct $i, j \in \{1, 2, 3\}$ are paired by Whitney disks in \mathcal{W} . The existence of such an order 1 non-repeating Whitney tower \mathcal{W} is equivalent to A having pairwise vanishing $\lambda_0(A_i, A_j) = 0 \in \mathbb{Z}[\pi_1 X]$ for $i \neq j$ (Exercise 3.9.2), which we can succinctly express as $\lambda_0(A) = 0$.

Recall from section 1.18.1 that if t_p associated to $p \in \mathcal{W}$ has univalent vertices labeled distinctly by 1, 2 and 3, then t_p is called a *non-repeating tree* and p is called a *non-repeating intersection*.

Denote by $\Lambda_1 := \Lambda_1(\pi_1 X)$ the quotient of the free abelian group on order 1 decorated non-repeating trees by the AS *antisymmetry* and HOL *holonomy* relations of Figure 3.4.

For $A \pitchfork X$ supporting an order 1 non-repeating Whitney tower \mathcal{W} , we define the *order 1 non-repeating intersection invariant* $\lambda_1(A)$ in the quotient of Λ_1 by the INT *intersection* relations shown in Figure 3.5:

$$\lambda_1(A) := \sum \epsilon_p \cdot t_p \in \Lambda_1 / \text{INT}$$




Figure 3.4. The *Antisymmetry* and *Holonomy* relations; $a, b, c, g \in \pi_1 X$.

where the sum is over all order 1 non-repeating intersections ρ in \mathcal{W} .




Figure 3.5. The INT *intersection* relations in the target of λ_1 , with $S : S^2 \looparrowright X$ varying over generators for $\pi_2(X)$ as a $\pi_1 X$ -module.

We next discuss how the AS *antisymmetry*, HOL *holonomy* and INT *intersection* relations in the target of λ_1 account for indeterminacies due to choices in the construction of the (based, oriented) order 1 non-repeating Whitney tower \mathcal{W} :

The AS relations (Figure 3.4 left) account for the choice of orientations on the Whitney disks via a fixed convention choice for the induced vertex orientation of the trees (section 1.19) so that signs of the order 1 intersections are well-defined (section 4.7). (In this non-repeating order 1 setting, AS relations could in fact be avoided by using the cyclic ordering of the distinct labels to prescribe orientations on all the Whitney disks in \mathcal{W} , cf. Exercise 3.9.8.)

The HOL relations (Figure 3.4 right) account for the choices of whiskers on the trivalent vertices, since changing such a whisker corresponds to simultaneous right multiplication on the three group elements decorating the edges of the tree. We remark that by the splitting operation (section 4.3) it can be arranged that each Whitney disk contains exactly one tree, so that in a split Whitney tower choosing a whisker for each trivalent vertex is the same as choosing a whisker for each Whitney disk.

The INT relations (Figure 3.5 left) account for choices of the interiors of Whitney disks in \mathcal{W} , and depend on A and $\pi_2 X$ via the order 0 intersection pairing λ_0 . To clarify notation: The terms of $\lambda_0(A_k, S)$ decorating the edge of the tree in the left of Figure 3.5 denote a linear combination of signed trees (section 4.1.2). Here S is a 2-sphere that has been tubed into the Whitney disk $W_{(i,j)}$, and $\lambda_0(A_k, S)$ is computed using a whisker for S given by the tube together with a whisker for the trivalent vertex of the tree (i, j) in $W_{(i,j)}$. Note that the illustration of S on the right side of Figure 3.5 is schematic, as suggested by the dotted subarc.

The following characterization of $\lambda_1(A_1, A_2, A_3)$ shows in particular that it does not depend on the choice of order 1 non-repeating Whitney tower:

Theorem 3.2. $\lambda_1(A_1, A_2, A_3)$ only depends on the homotopy classes of the A_i , and the following three statements are equivalent:

- (i) $\lambda_1(A_1, A_2, A_3)$ vanishes.
- (ii) $A_1 \cup A_2 \cup A_3$ admits an order 2 non-repeating Whitney tower.
- (iii) A_1, A_2, A_3 can be made pairwise disjoint by a homotopy.

Theorem 3.2 can be proved using arguments detailed in Section 4 (see Exercise 4.8.17), with the equivalence of statements (ii) and (iii) being a special case of Theorem 1.10 (see Exercise 3.9.6 for a proof in this case).

In the case that X is simply connected $\lambda_1(A_1, A_2, A_3)$ reduces to the Matsumoto triple [20] and the implication (i) \Rightarrow (iii) was shown by Yamasaki [35].

For more than three components one similarly defines the order 1 non-repeating intersection invariant $\lambda_1(A)$ of $A = A_1, A_2, \dots, A_m$ by counting, modulo the AS, HOL and INT relations, the distinctly-labeled decorated order 1 trees in the intersection forest of any order 1 non-repeating Whitney tower on A . Then $\lambda_1(A)$ is a homotopy invariant of A which vanishes if and only if A supports an order 2 non-repeating Whitney tower.

Note that the OR relations of section 3.2 are not needed in the order 1 setting because we are counting order 1 trees so edges can always be taken to be oriented towards the trivalent vertex.

3.6. Pulling apart m -tuples of spheres

As described in detail in [31], the above discussion gives the framework for a complete obstruction theory for *pulling apart m 2-spheres*, i.e. making the spheres pairwise disjoint by a homotopy: Theorem 1.10 says that the existence of an order $m - 1$ non-repeating Whitney tower on $A_1 \cup A_2 \cup \dots \cup A_m$ suffices to pull the A_i apart. And given an order n non-repeating Whitney tower \mathcal{W} supported by $A = A_1, A_2, \dots, A_m$, the intersection forest $t(\mathcal{W})$ represents an obstruction to the existence of an order $(n + 1)$ non-repeating Whitney tower. Denoting by $\Lambda_n(m)$ the abelian group generated by order n trees with distinctly labeled univalent vertices from $\{1, 2, \dots, m\}$ and edges decorated by elements of $\pi_1 X$, modulo OR, AS, HOL and IHX relations, where for $n \geq 2$ the IHX relation is as in 1.15 but with distinct univalent labels and with edge decorations, this obstruction lives in a quotient of $\Lambda_n(m)$ by order n *intersection relations* $\text{INT}_n(A)$ which correspond to changes in $t(\mathcal{W})$ coming from tubing Whitney disk interiors into 2-spheres. Since all the relations correspond to geometric manipulations of \mathcal{W} , the vanishing of the obstruction implies the existence of an order $(n + 1)$ non-repeating Whitney tower via the same order-raising maneuvers as in Section 4.6, which are all homogeneous in the univalent labels.

There are two main challenges here: One is concisely formulating the $\text{INT}_n(A)$ intersection relations, which for $n \geq 2$ can be non-linear. And the other is showing that no additional relations other than OR, AS, HOL, IHX and $\text{INT}_n(A)$ are needed to ensure that $\lambda_n(A)$ does not depend on the choice of \mathcal{W} . This second issue essentially amounts to showing that $\lambda_n(A)$ does not depend on the choices of Whitney disk boundaries. Evidence that the AS, HOL and IHX relations suffice to give independence of Whitney disk boundaries comes from Theorem 8 of [31] which states that $\lambda_n(A) \in \Lambda_n(m)$ does not depend on the choice of \mathcal{W} in the setting that the components A_i are properly immersed disks into the 4-ball, where $\text{INT}_n(A)$ relations are trivial (and OR is not relevant).

By [31, Lem.19], $\Lambda_n(m)$ is isomorphic to the direct sum of $\binom{m}{n+2} n!$ -many copies of the integral group ring $\mathbb{Z}[\pi_1 X^{(n+1)}]$ of the $(n + 1)$ -fold cartesian product $\pi_1 X^{(n+1)} = \pi_1 X \times \pi_1 X \times \dots \times \pi_1 X$. Note that $\Lambda_n(m)$ is trivial for $n \geq m - 1$ since an order n univalent tree has $n + 2$ univalent vertices. The absence of torsion in $\Lambda_n(m)$ is in alignment with the fact that the torsion subgroup (which is only 2-torsion) of the tree group \mathcal{T}_n of Definition 2.1 (repeating labels allowed) corresponds to the obstructions to converting twisted Whitney towers to framed Whitney towers in the 4-ball [9]. Such obstructions are not relevant in the non-repeating

setting since the boundary-twisting operation (Figure 2.2) can be used to eliminate twisted Whitney disks at the cost of only creating intersections whose trees do not have distinctly-labeled univalent vertices.

See [31] for further discussion of pulling apart m components, including examples where $\lambda_n(A)$ is well-defined.

The case of pulling apart 4-tuples of spheres in a simply connected 4-manifold is already interesting, and only partially understood. In [31, Sec.8] the order 2 invariant $\lambda_2(A)$ for $A = A_1, A_2, A_3, A_4$ supporting an order 2 non-repeating tower \mathcal{W} in a simply connected X is examined in detail, including a precise formulation of the $\text{INT}_2(A)$ relations which account for choices of Whitney disk interiors.

The following is Conjecture 29 of [31]:

Conjecture 3.3. *For any $A = A_1, A_2, A_3, A_4$ supporting an order 2 non-repeating tower \mathcal{W} in X , the components of A can be pulled apart if and only if $\lambda_2(A) \in \Lambda_2/\text{INT}_2(A)$ as defined in [31, Sec.8] vanishes.*

As explained in [31, Sec.8.3.6], the computation of the image of the $\text{INT}_2(A)$ relations in $\Lambda_2(4) \cong \mathbb{Z} \oplus \mathbb{Z}$ leads to some interesting number theory [18].

3.7. The order 1 self-intersection invariant τ_1

Having considered non-repeating order 1 invariants which generalize the classical order 0 intersection pairing and fit into an obstruction theory for pulling apart m -tuples of spheres, we next describe the order 1 generalization of the classical self-intersection invariant which fits into an obstruction theory for the existence of order n framed Whitney towers.

For simplicity we restrict attention to the case of a single immersed 2-sphere $A : S^2 \looparrowright X^4$ with $\tau_0(A) = 0$ (or in classical notation $\mu_0(A) = 0$). The vanishing of $\tau_0(A)$ means that A admits an order 1 framed Whitney tower \mathcal{W} , and we want to “count” the decorated trees t_p associated (as in section 3.4) to the order 1 intersections p in \mathcal{W} to define an order 1 invariant $\tau_1(A)$ which does not depend on the choice of \mathcal{W} and gives the complete obstruction to A bounding an order 2 framed Whitney tower.

Denote by $\tilde{\tau}_1 := \tilde{\tau}_1(\pi_1 X)$ the quotient of the free abelian group on decorated order 1 trees by the above AS, HOL relations (Figure 3.4) and also the new FR *framing relations* shown in Figure 3.6. The necessity of these FR relations is illustrated (schematically) on the right of Figure 3.6, which shows how performing opposite boundary-twists (Figure 2.2) on a framed Whitney disk W along different arcs of ∂W yields a new framed Whitney disk that has two interior intersections with A whose corresponding trees equal to the terms of an FR relation. Univalent labels are dropped from our trees since we are now working with just a single component. So the previously defined AS and HOL relations are now expressed in $\tilde{\tau}_1$ by the same equations but without univalent labels.




Figure 3.6. The FR framing relations; $a, b \in \pi_1 X$

We also need to refine the INT relations to account for the effect that tubing into a sphere can have on Whitney disk twistings. For instance, if a framed Whitney disk is tubed into a sphere S whose normal bundle has non-trivial 2nd Stiefel-Whitney number $\omega_2(S)$, then a framed Whitney disk can be recovered by boundary-twisting. Figure 3.7 shows these refined INT relations in the target of τ_1 , with the edge decoration $\omega_2(S)$ identified with 0 or 1 in $\mathbb{Z}[\pi_1 X]$ depending on whether it is trivial or not. This is explained in complete detail in section 4.1, where the full definition of the INT relations is given, which also allows for S to be an immersed \mathbb{RP}^2 in certain cases depending on the elements a, b .

INT:

Figure 3.7. The INT *intersection* relations in the target of τ_1 ; with $a, b \in \pi_1 X$, and $S : S^2 \looparrowright X$ varying over generators for $\pi_2(X)$ as a $\pi_1 X$ -module.

So for A with vanishing $\tau_0(A)$, choose an order 1 framed Whitney tower \mathcal{W} on A and define the *order 1 framed intersection invariant*:

$$\tau_1(A) := \sum \epsilon_p \cdot t_p \in \tilde{\mathcal{T}}_1/\text{INT}$$

where the sum is over all order 1 intersections p in \mathcal{W} .

The following result of [29] shows that $\tau_1(A)$ is an obstruction to homotoping A to an embedding, and in particular that it does not depend on the choice of \mathcal{W} :

Theorem 3.4. $\tau_1(A)$ only depends on the homotopy class of A , and the following four statements are equivalent:

- (i) $\tau_1(A)$ vanishes.
- (ii) A admits an order 2 framed Whitney tower (cf. Definition 1.7).
- (iii) A admits a height 1 Whitney tower (cf. Definition 1.12).
- (iv) A is stably homotopic to an embedding.

Here A being *stably homotopic to an embedding* means that A is homotopic to an embedding in the connected sum $X \#^n S^2 \times S^2$ of X with some number n copies of $S^2 \times S^2$.

See section 4.1 for proof that $\tau_1(A)$ is a well-defined homotopy invariant. The equivalence of statements (i) and (ii) is shown in section 4.2.

See [29, Thm.2] for the equivalence of statements (i) and (iii); and [28, Cor.1] for the equivalence of statements (i) and (iv).

3.8. Examples and questions

If X is simply-connected then the target $\tilde{\mathcal{T}}_1/\text{INT}$ of $\tau_1(A)$ is $\mathbb{Z}/2\mathbb{Z}$ or 0, depending on whether A is spherically characteristic or not. Here A is *spherically characteristic* if $\lambda_0(A, S) \equiv \lambda_0(S, S) \pmod{2}$ for all $S \in \pi_2 X$. In this simply-connected case $\tilde{\mathcal{T}}_1$ is generated by the order 1 tree with all trivial edge decorations, which is 2-torsion by the AS relations, and is equal to zero by the INT relations if A is not spherically characteristic.

For example, $3\mathbb{CP}^1 \looparrowright \mathbb{CP}^2$ is spherically characteristic, and Figure 3.8 shows the computation of $\tau_1(3\mathbb{CP}^1) = 1 \neq 0 \in \tilde{\mathcal{T}}_1/\text{INT}(3\mathbb{CP}^1) \cong \mathbb{Z}/2\mathbb{Z}$, so $3\mathbb{CP}^1$ is not homotopic to an embedding (compare [17]).




Figure 3.8. Left: The boundary circle in S^3 of the sum of three parallel copies of the $+1$ -framed 2-handle in a standard handle-body description of $\mathbb{C}P^2$. Moving to the right: The track into the 0-handle B^4 of a null-homotopy of this circle describes the rest of $3\mathbb{C}P^1$, with two self-intersections admitting the framed Whitney disk W that has a single transverse intersection with $3\mathbb{C}P^1$.

For X not necessarily simply-connected, taking the quotient of $\tilde{\mathcal{T}}_1$ by $\pi_1 X \rightarrow 1$ yields $\tau_1(A) \in \mathbb{Z}/2\mathbb{Z}$ or 0 depending on whether A is spherically characteristic or not, and in this quotient target τ_1 reduces to the spherical *Kervaire–Milnor* invariant km of [13, 33].

Even after trivializing all $\pi_1 X$ -decorations $\tau_1 = km$ sees global information in *closed* 4-manifolds:

Theorem 3.5 (Freedman–Kirby, Kervaire–Milnor, Stong). *Suppose X is a smooth closed 4-manifold, and $H_2(X; \mathbb{Z}/2\mathbb{Z})$ is spherical. If $A : S^2 \looparrowright X$ is characteristic and $\mu_0(A) = 0$, then*

$$(\pi_1 X \rightarrow 1) : \tau_1(A) \mapsto \frac{A \cdot A - \text{signature}(X)}{8} \pmod{2}$$

See the end of [33] for a proof.

On the other hand, for $\pi_1 X$ non-trivial the cardinality $|\tilde{\mathcal{T}}_1(\pi_1 X)|$ can be large. For example, if $\pi_1 X$ is left-orderable and the INT relations are trivial then $\tilde{\mathcal{T}}_1(\pi_1 X)$ is isomorphic to $\mathbb{Z}^\infty \oplus (\mathbb{Z}/2\mathbb{Z})^\infty$ (see [28, Prop.2.3.1]).

Let Γ be a finitely presented group, and let g be any element of $\tilde{\mathcal{T}}_1(\Gamma)$. Then one can find a 4-manifold X with non-empty boundary such that $\pi_1 X = \Gamma$, and $A : S^2 \looparrowright X$ with $\text{INT}(A)$ trivial and $\tau_1(A) = g$ (Exercise 3.9.9).

But finding examples of non-trivial $\tau_1(A)$ in *closed* 4-manifolds that depend on non-trivial edge decorations appears to be difficult, and no such examples are currently known. Let $(1, 1, 1)$ denote the order 1 tree with all three edges decorated by the trivial element $1 \in \pi_1 X$. Then $(1, 1, 1)$ is 2-torsion by AS (or FR) relations, and generates an order 2 subgroup of $\tilde{\mathcal{T}}_1(\pi_1 X)$, for any $\pi_1 X$. We have the following open “realization problem” for closed X :

Open Problem 3.6. *Find an example of $A \looparrowright X$ such that the following three statements all hold:*

- (i) X is closed, and
- (ii) $|\tilde{\mathcal{T}}_1(\pi_1 X)/\text{INT}(A)| \geq 3$, and
- (iii) $\tau(A)$ is not contained in the subgroup generated by $[(1, 1, 1)] \in \tilde{\mathcal{T}}_1(\pi_1 X)/\text{INT}(A)$.

Poincaré duality contributes to the difficulty of Problem 3.6 since if $\lambda_0(A, S) = 1$ for some $S \in \pi_2 X$, then $|\tilde{\mathcal{T}}_1(\pi_1 X)/\text{INT}(A)| \leq 2$ by the INT relations and τ_1 reduces to km .

An easier to state, but possibly more difficult, open realization problem is: Find an example of $A \looparrowright X$ such that $km(A) = 0$ but $\tau_1(A) \neq 0$.

In light of Theorem 3.5 and Problem 3.6 one is led to the question:

Question 3.7. *What global information is carried by non-trivial π_1 -decorations in τ_1 in closed 4-manifolds?*

In comparison with the classifications of order n framed and twisted Whitney towers on 2-disks in the 4-ball ([11] and Section 2), the higher-order invariants invariants for 2-spheres in 4-manifolds have barely begun to be defined, even in simply connected 4-manifolds. Defining an order 2 framed invariant is already an interesting challenge:

Open Problem 3.8. For $A : S^2 \looparrowright X$ with vanishing $\tau_1(A)$, formulate and prove invariance of an order 2 framed invariant $\tau_2(A)$ whose vanishing is equivalent to A admitting a framed order 3 Whitney tower.

In addition to the issues that arise as in Conjecture 3.3, here one has to also keep track of Whitney disk twistings.

3.9. Section 3 Exercises

3.9.1. Exercise:

Suppose A_i and A_j are oriented immersed 2-spheres in an oriented 4-manifold X , and p and q are oppositely signed intersections in $A_i \pitchfork A_j$ with equal associated group elements $g_p = g_q \in \pi_1 X$. Check that the union of a path in A_i from p to q with a path in A_j from q to p is a null-homotopic loop in X . Conclude that p and q admit an immersed Whitney disk.

3.9.2. Exercise:

For a collection of immersed 2-spheres $A = A_1, A_2, \dots, A_m \looparrowright X^4$, the vanishing of $\lambda_0(A)$ (section 3.2) is equivalent to the pairwise vanishing of the classical intersection pairing $\lambda_0(A_i, A_j) \in \mathbb{Z}[\pi_1 X]$ for $i \neq j$ (section 3.1). Show that $\lambda_0(A)$ vanishes if and only if A supports an order 1 non-repeating Whitney tower.

3.9.3. Exercise:

Show that any twisted Whitney disk can be converted to a framed Whitney disk having the same boundary by applying the boundary-twisting operation of section 2.2.

3.9.4. Exercise:

Show that $\tau_0(A)$ vanishes if A supports an order 1 framed Whitney tower.

3.9.5. Exercise:

Show that A supports an order 1 framed Whitney tower if $\tau_0(A)$ vanishes.

3.9.6. Exercise:

If A_1, A_2, A_3 support an order 2 non-repeating Whitney tower, show that A_1, A_2, A_3 can be made pairwise disjoint by a homotopy.

3.9.7. Exercise:

If order 1 intersections $p, q \in W_{(i,j)} \pitchfork A_k$ are paired by an order 2 Whitney disk $W_{((i,j),k)}$, then p and q have opposite signs by the definition of Whitney disks. Check that the order 1 decorated trees t_p and t_q associated to such p and q are equal for appropriate choices of trivalent whiskers.

3.9.8. Exercise:

Show that in the target of $\lambda_1(A_1, A_2, A_3)$ of section 3.5 the AS relations could be avoided by using the cyclic ordering of the distinct labels to prescribe orientations on all the Whitney disks in \mathcal{W} (via a choice of positive or negative corner convention cf. section 1.19).

3.9.9. Exercise:

Let Γ be a finitely presented group, and let t be any order 1 decorated tree representing an element $[t] \in \tilde{\mathcal{T}}_1(\Gamma)$. Find a 4-manifold X with non-empty boundary such that $\pi_1 X = \Gamma$, and $A : S^2 \looparrowright X$ with $\text{INT}(A)$ trivial and $\tau_1(A) = [t]$.

3.9.10. Exercise:

Generalizing the previous exercise, let Γ be a finitely presented group, and let g be any element of $\tilde{\mathcal{T}}_1(\Gamma)$. Find a 4-manifold X with non-empty boundary such that $\pi_1 X = \Gamma$, and $A : S^2 \looparrowright X$ with $\text{INT}(A)$ trivial and $\tau_1(A) = g$.

3.9.11. Exercise:

Generalizing the previous exercise, let Γ be a finitely presented group, let g be any element of $\tilde{\mathcal{T}}_1(\Gamma)$, and let $\{z_1, z_2, \dots, z_n\}$ be any collection $z_i \in \mathbb{Z}[\Gamma]$. Find a 4-manifold X with non-empty boundary such that $\pi_1 X = \Gamma$, and $A : S^2 \looparrowright X$ with $\text{INT}(A)$ determined by $\lambda_0(A, S_i) = z_i$ for S_i generating $\pi_2 X$, and $\tau_1(A) = g$.

4. Appendix

Section 4.1 gives a detailed proof of the homotopy invariance of the order 1 invariant $\tau_1(A)$ from section 3.7, and section 4.2 proves that $\tau_1(A)$ is the complete obstruction to A supporting an order 2 framed Whitney tower. This proves the first two statements of Theorem 3.4, and simpler versions of the analogous arguments prove the first two statements of Theorem 3.2 (Exercise 4.8.17).

The splitting of Whitney towers illustrated in Figure 1.10 is extended to twisted Whitney towers in section 4.3.

Whitney move versions of the IHX and twisted IHX relations are described in detail in section 4.5. These constructions are essential to the framed, twisted, and non-repeating order-raising obstruction theories. A proof of the twisted order-raising Theorem 2.5 is sketched in section 4.6.

4.1. Homotopy invariance of $\tau_1(A)$

Here we show that for $A : S^2 \looparrowright X^4$ the order 1 framed intersection invariant $\tau_1(A)$ of Theorem 3.4 in section 3.7 only depends on the homotopy class of A .

The definition of $\tau_1(A)$ requires that A supports an order 1 framed Whitney tower \mathcal{W} , the existence of which is equivalent to A having vanishing order 0 self-intersection invariant $\tau_0(A) = 0$ by Theorem 3.1. As recalled and clarified below, $\tau_1(A)$ is determined by the intersection forest $t(\mathcal{W})$, which in this setting is a multiset of signed decorated order 1 trees, one for each transverse intersection between A and a Whitney disk in \mathcal{W} . The main challenge is to show that the element represented by $t(\mathcal{W})$ in the target of $\tau_1(A)$ does not depend on the choices of oriented Whitney disks in \mathcal{W} . Then invariance under homotopies of A will be shown in section 4.1.9.

4.1.1. Order 1 tree notation

We introduce the following streamlined notation for the order 1 decorated trees of section 3.4 in our current 1-component setting: Since now the univalent labels are not relevant, a tree is determined by a cyclic ordering of the elements of $\pi_1 X$ decorating the edges. By appropriately choosing the whisker on the trivalent vertex, one edge decoration on each tree can be normalized to the trivial element $1 \in \pi_1 X$. So for $a, b \in \pi_1 X$ we will use the notation (a, b) to denote the decorated order 1 tree with one edge decorated by 1 and the other two edges decorated by a and b respectively, with the vertex orientation given by the cyclic ordering $1 \rightarrow a \rightarrow b$.

As before, any edge decorations in $\mathbb{Z}[\pi_1 X]$ are understood to be linear combinations of trees.

4.1.2. Relations in target of τ_1

Recall from section 3.7 that $\tilde{\mathcal{T}}_1 := \tilde{\mathcal{T}}_1(\pi_1 X)$ denotes the abelian group on decorated order 1 trees modulo the AS *antisymmetry*, HOL *holonomy* and FR *framing* relations (see just below), and that $\tau_1(A)$ takes values in $\tilde{\mathcal{T}}_1$ modulo the INT *intersection* relations. Since these last relations depend on A we sometimes denote them by INT(A) for emphasis.

The relations in the target of $\tau_1(A)$ expressed in $\mathbb{Z}[\pi_1 X] \times \mathbb{Z}[\pi_1 X]$ via the streamlined notation are:

$$\begin{aligned} \text{AS:} \quad & (a, b) = -(b, a) \\ \text{HOL:} \quad & (a, b) = (ba^{-1}, a^{-1}) = (b^{-1}, ab^{-1}) \\ \text{FR:} \quad & (1, a) + (a, a) = 0 \\ \text{INT:} \quad & (a, \lambda_0(A, S) + \omega_2(S) \cdot 1) = 0 \end{aligned}$$

Here $a, b \in \pi_1 X$, and S ranges over generators for $\pi_2 X$ (as a module over $\pi_1 X$).

The AS and HOL relations imply that the decorated trees $(a, 1)$ and (a, a) are 2-torsion elements, which by FR are equal. This will be used in several places below.

Terms are expanded linearly in the INT relations, so

$$(a, \lambda_0(A, S) + \omega_2(S) \cdot 1) = (a, \lambda_0(A, S)) + (a, \omega_2(S) \cdot 1),$$

and

$$(a, \lambda_0(A, S)) = (a, \sum \epsilon_p \cdot g_p) := \sum \epsilon_p \cdot (a, g_p)$$

for $\lambda_0(A, S) = \sum_{p \in \text{Int} S} \epsilon_p \cdot g_p$.

For the case $a^2 = 1 \in \pi_1 X$ in the INT relations, we allow S to be an immersed \mathbb{RP}^2 representing a , i.e. $a \in \pi_1 X$ is the image of $\pi_1 \mathbb{RP}^2$. When S is an \mathbb{RP}^2 the order 0 intersection pairing $\lambda_0(A, S)$ is only well-defined up to right multiplication by a and change of sign, but the AS and HOL relations make $(a, \lambda_0(A, S))$ well-defined in $\tilde{\mathcal{T}}_1$. This will be discussed in detail in section 4.1.7 where this form of the INT relation is used in the proof that $\tau_1(A)$ is well defined.

4.1.3. Definition of $\tau_1(A)$

For A satisfying $\tau_0(A) = 0$, choose any order 1 framed Whitney tower \mathcal{W} on A and define:

$$\tau_1(A) := \sum \epsilon_p \cdot t_p \in \tilde{\mathcal{T}}_1 / \text{INT}(A)$$

where the sum is over all transverse intersections p between A and the Whitney disks in \mathcal{W} . We use the *positive corner* convention from section 1.19 for inducing cyclic orientations at the trivalent vertices from the Whitney disk orientations.

We remark that by eliminating all order 2 intersections using the pushing-down operation (Exercise 1.22.13) it may be arranged that the intersection forest $t(\mathcal{W})$ (section 1.13) is

equal to the sum defining $\tau_1(A)$, i.e. $\tau_1(A) = [t(\mathcal{W})] \in \tilde{\mathcal{T}}_1/\text{INT}(A)$. In the following this will be assumed for clarity, although it is not necessary.

4.1.4. Independence of Whitney disk interiors:

Let W be a Whitney disk in \mathcal{W} . If V is another framed Whitney disk with same boundary γ as W , then the union of W and V is a map of a 2-sphere, which might not be smooth along γ . We want to use this sphere to express, via an INT relation, the change in $\tau_1(A)$ due to replacing W by V . It is not obvious what intersections with A might be created by perturbing the sphere to be smooth and transverse to A , so we will proceed carefully.

Denote by α and β the two closed subarcs of $\alpha \cup \beta = \gamma := \partial W = \partial V$ that run between the self-intersections p and q of A that are paired by W and V .

A small collar of α in W determines a 1-dimensional subbundle w_α of $\nu_X A|_\alpha$, the normal D^2 -bundle to A restricted to α . Similarly, a small collar of α in V determines a 1-dimensional subbundle v_α of $\nu_X A|_\alpha$. Note that at each of the endpoints p and q of α there is a canonical 1-dimensional subspace of $\nu_X A|_\alpha$ given by a short subarc of β . We may assume that $w_\alpha = v_\alpha$ near p and q by a small isotopy, so there is a relative rotation number $m \in \mathbb{Z}$ of v_α with respect to w_α . A boundary-twist (Figure 2.2) on V changes m by ± 1 , so by performing m -many boundary-twists on V we can get a new Whitney disk V' such that the corresponding 1-dimensional subbundle v'_α of $\nu_X A|_\alpha$ has zero relative rotation number with respect to w_α , and hence it can be arranged by a small isotopy of V' near α that $v'_\alpha = w_\alpha$ everywhere along α .

Applying the same discussion and boundary-twisting n times along β , we can assume that we have also arranged that a small collar v'_β of β in V' satisfies $v'_\beta = w_\beta \subset \nu_X A|_\beta$ everywhere along β .

Now the unions $w := w_\alpha \cup w_\beta$ and $v' := v'_\alpha \cup v'_\beta$ are collars of γ in W and V' which are equal $w = v'$.

Removing these collars and gluing we get a map of a 2-sphere $S := (W \setminus w) \cup (V' \setminus v')$ which is smooth except along the gluing circle $\gamma' := \partial(W \setminus w) = \partial(V' \setminus v')$. Since γ' is disjoint from A , perturbing S to be smoothly immersed will not create any new intersections with A . (Note that we do not need to control self-intersections of S because they do not correspond to contributions to τ_1 .) The intersections between S and A are transverse and consist of $A \pitchfork W$ and $A \pitchfork V'$, where $A \pitchfork V'$ consists of $A \pitchfork V$ together with the $(m+n)$ -many intersections of $A \pitchfork V'$ which were created by the boundary-twists on V . These boundary-twist intersections correspond to the trees $m(1, a) + n(a, a)$, where m twists were done along the arc $\alpha \subset \gamma$ whose corresponding edge decoration has been normalized to 1, and n twists were done along the other arc $\beta \subset \gamma$ whose corresponding edge decoration is a (Figure 3.6, with $b = 1$). Taking the orientation of S to be induced by the orientation of V' and the opposite orientation of W , we have that the difference in $\tau_1(A)$ due to replacing W by V is $(a, \lambda_0(A, S)) + m(1, a) + n(a, a)$. It just remains to check that $\omega_2(S) \equiv m+n \pmod{2}$, so that this difference is an INT relation.

To compute $\omega_2(S)$ as a mod 2 self-intersection number, observe that S is homotopic to $S' = W \cup V'$ by a homotopy supported near the common collars $w = v'$ that were deleted to form S (Exercise 4.8.1). Now we can assume that S' has been perturbed to be smooth along γ without worrying about creating intersections with A since we will only be counting the mod 2 intersections between S' and a normal push-off S'' . Form S'' by taking any Whitney parallel push-off of γ and extending over W and V' by Whitney sections to get the two hemispheres of S'' , and count the number of intersections in $S' \pitchfork S''$: The number of intersections between W and its push-off must be even, since W is framed. The number of intersections between V' and its push-off must equal $m+n \pmod{2}$, since V' was created by $m+n$ boundary-twists on the framed V . Each transverse intersection between W and V' will contribute two intersections between their push-offs.

So we have $\omega_2(S) \equiv m + n \pmod{2}$, and it follows that the change in $\tau_1(\mathcal{W})$ resulting from replacing W by V is exactly described by the INT(A) relation:

$$(a, \lambda_0(A, S)) + m(1, a) + n(a, a) = (a, \lambda_0(A, S) + \omega_2(S) \cdot 1)$$

4.1.5. Independence of Whitney disk boundaries:

To show that $\tau_1(\mathcal{W})$ does not depend on the choices of boundaries of the Whitney disks, for fixed pairings of the self-intersections in A , it is convenient to temporarily weaken the definition of an order 1 framed Whitney tower by allowing transverse intersections among the boundaries of the Whitney disks in \mathcal{W} (following [13, Sec.10.8] and [29]). The definition of τ_1 is extended to such Whitney towers by assigning trees to the boundary intersections between Whitney disks in the following way.

The Whitney disk orientations induce orientations on their boundaries via the usual convention that $\overrightarrow{\partial W}$ together with a second inward pointing vector give the orientation of W . We will use the notation $\partial_+ W$ to indicate the boundary arc of W that is oriented towards the positive self-intersection of A paired by W , and $\partial_- W$ for the boundary arc of W that is oriented towards the negative self-intersection paired by W .

Let $p \in \partial_\epsilon W \cap \partial_\delta V$, for $\epsilon, \delta \in \{+, -\}$, be a point such that the ordered pair of tangent vectors $(\overrightarrow{\partial_\epsilon W}, \overrightarrow{\partial_\delta V})_p$ is equal to the orientation of A at p . Choose whiskers on the trivalent vertices of the rooted trees associated to W and V so that the edges dual to $\partial_\epsilon W$ and $\partial_\delta V$ are each decorated by the trivial element $1 \in \pi_1 X$. This determines elements a and b decorating the other edges of the two respective trees. See Figure 4.1. Define the signed tree $\epsilon_p \cdot t_p$ associated to such a p by:

$$(4.1) \quad \epsilon_p \cdot t_p := -\epsilon\delta \cdot (a^\epsilon, b^\delta)$$

where $\epsilon_p, \epsilon, \delta \in \{+, -\} \cong \{+1, -1\}$. The reason for the minus sign in Equation (4.1) will be made clear in Figure 4.3 below.




Figure 4.1. The case $\epsilon = + = \delta$ in Equation (4.1): The illustrated intersection point $p \in \partial_\epsilon W \cap \partial_\delta V$ between Whitney disk boundaries is assigned the signed tree $\epsilon_p \cdot t_p = -\epsilon\delta(a^\epsilon, b^\delta) = -(a, b)$.

One can check that this definition of t_p does not depend on the choices made (Exercise 4.8.2). The extended version of $\tau_1(A)$ is defined by including all such t_p in the sum. Since all boundary intersections can be eliminated by finger moves which create interior intersections having the exact same trees (Figure 4.2), this extended definition can always be reduced to the original one.

Properly interpreted, the formula assigning t_p to $p \in \partial W \cap \partial V$ also works when $W = V$. For instance, for $p \in \partial_- W \cap \partial_+ W$ such that the orientation of A is equal to $(\overrightarrow{\partial_- W}, \overrightarrow{\partial_+ W})_p$, then




Figure 4.2. Eliminating $p \in \partial_+ W \cap \partial_+ V$ from Figure 4.1 by pushing across the positive self-intersection paired by W creates a positive intersection $p' \in A \pitchfork V$ with signed tree $\epsilon_{p'} \cdot t_{p'} = +(b, a) = -(a, b) = \epsilon_p \cdot t_p$.

$\epsilon_p \cdot t_p = -(a, a^{-1})$, where the rooted tree associated to W has decorations 1 and a on the edges dual to the $\partial_+ W$ and $\partial_- W$ boundary arcs respectively. (Exercise 4.8.2.)

The proof of independence of Whitney disk boundaries now goes as follows. For a fixed choice of pairings of self-intersections induced by a given collection of Whitney disks, any other configuration of Whitney disk boundaries can be achieved by a regular homotopy of (collars of) the given Whitney disk boundaries, fixing the self-intersection points of A (Clarification: we mean here that this regular homotopy is induced by a regular homotopy of the preimages of the Whitney disk boundaries in the domain of A , and extends to a regular homotopy of collars of the Whitney disks in X). During such a homotopy, τ_1 does not change since boundary intersections come and go in canceling pairs, or accompanied by a canceling interior intersection when pushing over a self-intersection point of A , see Figure 4.3 for one case and Exercise 4.8.3 for the others. (This step uses the fact that the domain of A is simply connected.)




Figure 4.3. Pushing $\partial_+ W$ into $\partial_+ V$ across the positive self-intersection of A paired by V creates $p \in \partial_+ W \cap \partial_+ V$ and $q \in A \pitchfork W$ with algebraically canceling signed trees $\epsilon_p \cdot t_p = -(a, b) = -\epsilon_q \cdot t_q$.

4.1.6. Independence of pairings of self-intersections:

Let W be a Whitney disk in \mathcal{W} pairing self-intersections p and q , and let W' be another Whitney disk in \mathcal{W} pairing self-intersections p' and q' , such that all these self-intersections

determine the same element of $\pi_1 X$. Then there exists a framed Whitney disk V' pairing p' and q such that $\partial V'$ is disjoint from all singularities in \mathcal{W} (including Whitney disk boundaries, cf. Figure 4.2). Now a framed Whitney disk V pairing p and q' can be constructed from W and W' (minus small corners) and a parallel of V' (with the orientation reversed) as illustrated in Figure 4.4.

Replacing W and W' with V and V' in \mathcal{W} does not change $\tau_1(A)$ since the order 1 decorated trees corresponding to the intersections $A \pitchfork V'$ cancel with those corresponding to the oppositely-signed parallel copy of V' in V (Exercise 4.8.4). Any choice of pairings can be achieved by iterating this construction.




Figure 4.4. Transverse intersections between Whitney disks and A are not shown. This picture is otherwise accurate up to diffeomorphism since W , W' and V' can be assumed to be framed and embedded. The dotted sub-arc of the doublepoint loop for p is understood to extend outside the 4-ball of the figure.

4.1.7. The INT relation for $\mathbb{R}P^2$ s

For use in the next step of the proof, we clarify here the INT relation $(a, \lambda_0(A, R) + \omega_2(R) \cdot 1) = 0 \in \tilde{\mathcal{T}}_1$ for the case where $a^2 = 1$ and $R : \mathbb{R}P^2 \looparrowright X$.

The pairing $\lambda_0(A, R)$ sums the signed group elements $\epsilon_p \cdot g_p$ associated to each $p \in A \pitchfork R$ as in section 3.1, except that because $\mathbb{R}P^2$ is neither simply connected nor orientable now both ϵ_p and g_p depend on the choice of sheet-changing loop through p . Let d be a sheet-changing loop through p , and let d_R be the path $d \cap R$. Then $g_p = [d] \in \pi_1 X$ is the group element associated to p using d , and by definition the sign ϵ_p is gotten by transporting a local orientation of R from the basepoint of R back along d_R to p , where it is paired with the orientation of A at p for comparison with the orientation of X . Any different choice of d_R changes g_p by right multiplication by a^n and changes ϵ_p by multiplication by $(-1)^n$ (Exercise 4.8.5), so $\lambda_0(A, R) \in \mathbb{Z}[\pi_1 X]$ is only well-defined up to the relations $g = -ga$, for all $g \in \pi_1 X$.

In the setting of the INT relations, the local orientation of R comes from the orientation of the Whitney disk that R has been tubed into, and we have $(a, g) = -(a, ga)$ by the HOL and AS relations. So $(a, \lambda_0(A, R))$ is well-defined in $\tilde{\mathcal{T}}_1$.

As for 2-spheres, the second Stiefel–Whitney number $\omega_2(R) \in \mathbb{Z}_2$ is computed as $|R \pitchfork R'| \bmod 2$, where R' is a parallel copy of R . (Recall that $(a, \omega_2(R) \cdot 1)$ is 2-torsion in $\tilde{\mathcal{T}}_1$.)

4.1.8. Independence of sheet choices:

For $A : S^2 \looparrowright X$, let p and q be a positive and a negative transverse self-intersection of A , and denote the preimages by $A^{-1}(p) = \{x, x'\} \subset S^2$, and $A^{-1}(q) = \{y, y'\} \subset S^2$.

If p and q have common group element a , then any Whitney disk W pairing p and q induces a pairing of $\{x, x'\}$ with $\{y, y'\}$ since each arc of ∂W runs between a sheet of A around p and a sheet of A around q .

Framed Whitney disks exist for *both* of the two pairing choices $x \leftrightarrow y, x' \leftrightarrow y'$ and $x \leftrightarrow y', x' \leftrightarrow y$ if and only if $a^2 = 1 \in \pi_1 X$ (Exercise 4.8.6). So we need to show that for any a such that $a^2 = 1$, $\tau_1(A)$ does not depend on these *choices of preimage pairings*, also called *choices of sheets* ([32, Sec.4]).




Figure 4.5. In S^2 : The quadrilateral $r = c_+ \cup d_- \cup c_- \cup d_+$ bounds the 2-cell D .

Let W and V be framed Whitney disks corresponding to the two ways of pairing the preimages of such a pair p and q having group element a with $a^2 = 1$. The union of the preimages of the Whitney disk boundary arcs $A^{-1}(\partial W) \cup A^{-1}(\partial V)$ is a quadrilateral r in S^2 which is the union of the two pairs of arcs $c_{\pm} := A^{-1}(\partial_{\pm} W)$ and $d_{\pm} := A^{-1}(\partial_{\pm} V)$ (see Figure 4.5). Since $\tau_1(A)$ does not depend on the choices of boundaries of W and V , we may arrange that r is embedded and bounds a 2-cell D in S^2 such that A restricts to an embedding on D . (For later use we have chosen ∂V to coincide with ∂W near p and q , as indicated in Figure 4.5.)

The union $A(D) \cup W \cup V$ defines the image of a map $R : \mathbb{RP}^2 \rightarrow X$ representing a (Exercise 4.8.7). This R contains the transverse intersections $A \pitchfork W$ and $A \pitchfork V$, but R is not transverse to A along $A(D)$, and R is not smooth along $\rho := A(r)$. We will describe a perturbation of R to a smoothly immersed $R' : \mathbb{RP}^2 \looparrowright X$ transverse to A such that the perturbation creates exactly new intersections $R' \pitchfork A$ which correspond to $\omega_2(R')$. Then the change in $\tau_1(A)$ due to replacing W by V will be given by $(a, \lambda_0(A, R') + \omega_2(R') \cdot 1)$.

The perturbation of R to R' uses the following lemma, which will be proved below:

Lemma 4.1. *There exists a push-off ρ' of ρ such that ρ' is normal to A and ρ' restricts to Whitney parallels of ∂W and ∂V .*

See Remark 1.6 for an explanation of “Whitney parallel”.

Given ρ' as in Lemma 4.1, we define $R' : \mathbb{RP}^2 \looparrowright X$ as follows. Extend ρ' across W and V to get Whitney parallels W' and V' . Extend ρ' generically across $A(D)$ to get $A(D)'$ which is transverse to $A(D)$.

This defines the image $W' \cup V' \cup A(D)'$ of a map $\mathbb{RP}^2 \rightarrow X$. Now smooth the corners of $W' \cup V' \cup A(D)'$ near ρ' by a perturbation supported away from A to get $R' : \mathbb{RP}^2 \looparrowright X$.

We have $A \pitchfork R' = (A \pitchfork W') \cup (A \pitchfork V') \cup (A \pitchfork A(D)')$.

By construction, the intersections $A \pitchfork W'$ are all parallel to the intersections $A \pitchfork W$, and the intersections $A \pitchfork V'$ are all parallel to the intersections $A \pitchfork V$. Also $A \pitchfork A(D)' = A(D) \pitchfork A(D)'$, since A restricts to an embedding on D .

We claim that $|A(D) \pitchfork A(D)'| \equiv \omega_2(R') \pmod{2}$. To check this claim, we can use a smooth perturbation of the topological $R : \mathbb{RP}^2 \rightarrow X$ above as a parallel of R' , and compute $\omega_2(R') = |R \pitchfork R'| \pmod{2}$, as in section 4.1.7. Recall that $R = A(D) \cup W \cup V$, and since we now are not concerned with controlling intersections with A we can take any small perturbation near ρ to make R smooth.

Since W and V are framed, and R' restricts to Whitney parallels of ∂W and ∂V , we have $W \pitchfork W' \equiv 0 \equiv V \pitchfork V' \pmod{2}$. So the only possible contributions to $|R \pitchfork R'| \pmod{2}$ come from $A(D) \pitchfork A(D)'$ as claimed.

The group element associated to any point in $A(D) \pitchfork A(D)' = A \pitchfork A(D)' \subset A \pitchfork R'$ is 1 or a , so the change in $\tau_1(A)$ due to replacing W by V is $(a, \lambda_0(A, R') + \omega_2(R') \cdot 1)$, since $(a, 1) = (a, a)$ is 2-torsion.

To show independence of sheet choice it just remains to prove Lemma 4.1.

Proof of Lemma 4.1. We want to define a push-off ρ' of $\rho = \partial_+ W \cup \partial_- V \cup -\partial_- W \cup -\partial_+ V$ which is normal to A and restricts to Whitney parallels of ∂W and ∂V .

Consider the normal disk-bundle $\nu_X W|_{\partial W}$ which is the restriction to ∂W of the normal disk-bundle $\nu_X W$ of W . Denote by $A_{\pm} \subset \nu_X W|_{\partial W}$ the sheets corresponding to A with $\partial_- W \subset A_-$ and $\partial_+ W \subset A_+$. Figure 4.6 shows an embedding of $\nu_X W|_{\partial W} \cong S^1 \times D^2$ into 3-space, as in Figure 1.6.

At p and at q , the two sheets A_{\pm} split each of $\nu_X W|_{\partial W}(q)$ and $\nu_X W|_{\partial W}(p)$ into four quadrants. We define a preferred quadrant at p and at q as follows. The orientation of A induces orientations of the A_{\pm} , which in turn induce orientations of ∂A_- and ∂A_+ via the usual convention. At the positive self-intersection p the preferred quadrant is bounded by the vectors $\overrightarrow{\partial A_+}$ and $\overrightarrow{\partial A_-}$, and at the negative self-intersection q the preferred quadrant is bounded by the vectors $-\overrightarrow{\partial A_+}$ and $-\overrightarrow{\partial A_-}$.

Define $\rho'(p)$ and $\rho'(q)$ by choosing a vector in each of these preferred quadrants (see left of Figure 4.6).

Extend $\rho'(p)$ and $\rho'(q)$ to a section $\rho'(\partial W)$ of $\nu_X W|_{\partial W}$ over ∂W such that $\rho'(\partial W)$ is normal to both sheets of A along ∂W , as in the right of Figure 4.6. Note that this section exists by our choices of preferred quadrants.

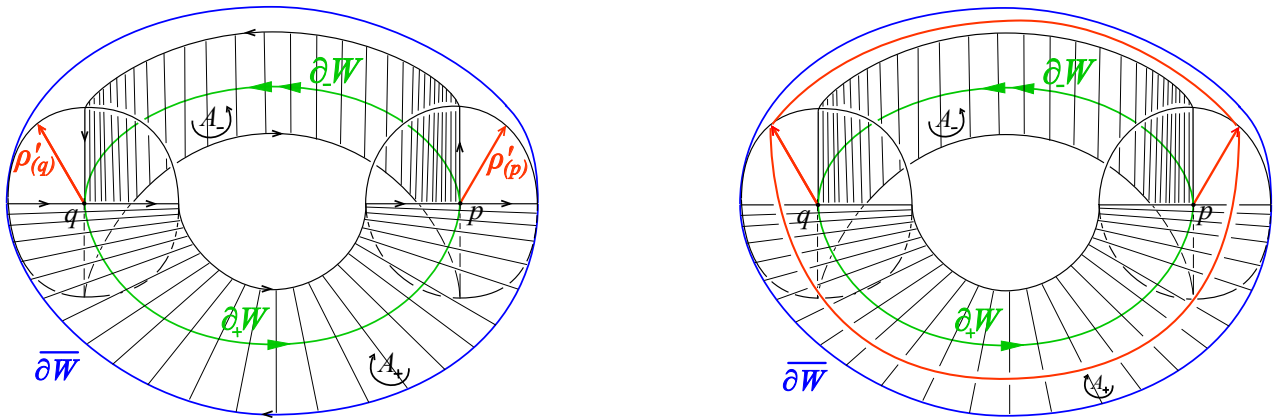


Figure 4.6. In $\nu_X W|_{\partial W}$: On the left, $\rho'(q)$ and $\rho'(p)$ in the preferred quadrants, and on the right, extended to a red Whitney section $\rho'(\partial W)$ which is normal to A . The blue ‘standard’ Whitney section $\overline{\partial W}$ is shown as a reference with Figure 1.6.

To define ρ' over ∂V , note that we may assume that $V = W$ near both p and q (Exercise 4.8.8), so $\nu_X V|_{\partial V}$ coincides with $\nu_X W|_{\partial W}$ near p and q . This means that the preferred quadrants also coincide, and the previously defined $\rho'(p)$ and $\rho'(q)$ can be extended along ∂V to a section $\rho'(\partial V)$ which is normal to the sheets of A .

So along $\rho = \partial_+ W \cup \partial_- V \cup -\partial_- W \cup -\partial_+ V$ we have the push-off ρ' assembled from the push-offs along the given oriented subintervals starting at $\rho(q)$: The section $\rho'(\partial_+ W)$ ending at $\rho'(p)$ which is the restriction to $\partial_+ W$ of $\rho'(\partial W)$, followed by the section $\rho'(\partial_- V)$ which is the restriction to $\partial_- V$ of $\rho'(\partial V)$ running from $\rho'(p)$ back to $\rho'(q)$, followed by the section $\rho'(-\partial_- W)$ which is the restriction to $\partial_- W$ of $\rho'(\partial W)$ but running from $\rho'(q)$ back to $\rho'(p)$, followed by the section $\rho'(-\partial_+ V)$ which is the restriction to $\partial_+ V$ of $\rho'(\partial V)$ but running from $\rho'(p)$ back to $\rho'(q)$. \square

We have so far shown that $\tau_1(A) \in \tilde{\tau}_1/\text{INT}(A)$ is independent of the choice of order 1 framed Whitney tower for a fixed immersion A .

4.1.9. Invariance of $\tau_1(A)$ under homotopies of A

It suffices to check invariance under regular homotopy by [25, Thm.1.2]. Suppose that $\tau_0(A)$ vanishes, and A is regularly homotopic to A' . By the homotopy invariance of τ_0 it follows that A' also supports an order 1 framed Whitney tower. Recall that up to isotopy, any generic regular homotopy from A to A' can be realized as a sequence of finitely many finger moves followed by finitely many Whitney moves.

An isotopy from A to A' extends to any Whitney tower \mathcal{W} on A to yield a Whitney tower \mathcal{W}' on A' with identical intersection forests $t(\mathcal{W}) = t(\mathcal{W}')$, so τ_1 is invariant under isotopy.

Since any Whitney move has a finger move as an “inverse”, there exists A'' which differs from each of A and A' by only finger moves (up to isotopy). Since a finger move is supported near an arc, it can be made disjoint from the Whitney disks in any pre-existing Whitney tower by a small isotopy, and the pair of intersections created by a finger move admit a local clean Whitney disk disjoint from any other Whitney disks. So any Whitney tower on A or A' gives rise to a Whitney tower on A'' yielding $\tau_1(A) = \tau_1(A'') = \tau_1(A')$, since τ_1 does not depend on the choice of Whitney tower on A'' . Here we are also using that $\text{INT}(A) = \text{INT}(A') = \text{INT}(A'')$, since A , A' and A'' are all homotopic.

4.2. $\tau_1(A)$ vanishes if and only if A supports an order 2 framed Whitney tower

Here we show the equivalence of statements (i) and (ii) in Theorem 3.4.

If A supports an order 2 framed Whitney tower \mathcal{V} , then $\tau_1(A)$ vanishes since \mathcal{V} is also an order 1 framed Whitney tower, and \mathcal{V} contains no unpaired order 1 intersections.

So assume that $\tau_1(A)$ vanishes, and let \mathcal{W} be an order 1 framed Whitney tower on A . The vanishing of $\tau_1(A) = [t(\mathcal{W})] \in \widetilde{\mathcal{T}}_1/\text{INT}$ means that the intersection forest $t(\mathcal{W}) = \sum \epsilon_p \cdot t_p$ lies in the span of the AS, HOL, FR and INT relators. Here we are considering $t(\mathcal{W})$ as a word in the span of decorated order 1 trees.

The construction of an order 2 framed Whitney tower on A will involve two main steps: First, geometric realizations of the relators will be used to modify \mathcal{W} so that $t(\mathcal{W})$ consists of pairs $\pm t_p$ of oppositely-signed isomorphic trees corresponding to *algebraically canceling pairs* of intersections. Then, further controlled modifications of \mathcal{W} will convert these algebraically canceling pairs into *geometrically canceling pairs* of intersections admitting Whitney disks. The issue here is that an algebraically canceling pair of intersections may lie in different Whitney disks (Figure 4.7), so achieving geometric cancellation will require “transferring” intersections from one Whitney disk to another.

In the following \mathcal{W} will not be renamed as controlled modifications are made.

4.2.1.

Towards algebraic cancellation: Perform a finger move on A guided by a circle representing any $a \in \pi_1 X$, and then tube the resulting clean local framed Whitney disk into a 2-sphere S . The resulting Whitney disk W has twisting $\omega(W) \equiv \omega_2(S) \pmod{2}$. If $\omega_2(S) = 0$, then after performing some interior twists on W it can be arranged that $\omega(W) = 0 \in \mathbb{Z}$. If $\omega_2(S) = 1$, then after performing one boundary-twist and some interior twists on W it can be arranged that $\omega(W) = 0 \in \mathbb{Z}$. So after recovering the framing on W it follows that $[t(\mathcal{W})] \in \widetilde{\mathcal{T}}_1$ is changed exactly by adding the INT relator $(\alpha, \lambda_0(A, S) + \omega_2(S) \cdot 1)$.

In the case where $a^2 = 1$, a similar but more complicated procedure changes $[t(\mathcal{W})] \in \widetilde{\mathcal{T}}_1$ exactly by adding $(\alpha, \lambda_0(A, R) + \omega_2(R) \cdot 1)$, where $R : \mathbb{R}P^2 \looparrowright X$ represents a . This will be shown below in Lemma 4.2.

By realizing INT relations in these ways we can arrange that $t(\mathcal{W})$ lies in the span of AS, HOL and FR relators. To realize any FR relator $(1, a) + (a, a)$ perform a finger move guided by a circle representing a to get a clean local framed Whitney disk W . Then performing two

opposite boundary twists around each boundary arc of W changes $t(\mathcal{W})$ exactly by $(1, a) + (a, a)$.

So by realizing INT and FR relators we can arrange that $t(\mathcal{W})$ lies in the span of AS and HOL relators. By re-choosing whiskers on the trivalent vertices of trees to realize HOL relations we can further arrange that $t(\mathcal{W})$ lies in the span of AS relators.

At this point it is convenient to *split* \mathcal{W} so that each Whitney disk in \mathcal{W} contains exactly one interior intersection with A . This splitting is accomplished by finger moves guided by arcs in the Whitney disks running from one boundary arc to the other, and is the easiest case of Lemma 4.3 (see also Figure 1.10). Now any tree $(a, b) \in t(\mathcal{W})$ can be changed to $-(a, b) \in t(\mathcal{W})$ simply by re-choosing the orientation of the corresponding Whitney disk.

As a result of these constructions we can assume that now the un-paired intersections in \mathcal{W} consist of algebraically canceling pairs.




Figure 4.7

4.2.2.

Towards geometric cancellation: Consider an algebraically canceling pair of intersections $p = A \# W$ and $q = A \# W'$, with $\epsilon \cdot t_p = (a, b)$ and $\epsilon \cdot t_q = -(a, b)$, where the b -decorated edge changes sheets at p and q respectively. See Figure 4.7, where the short dashed sub-arcs indicate where sheets extend outside the 4-ball shown in the figure.

We will describe a controlled modification of \mathcal{W} that “transfers” p over to W' so that p and q admit a framed order 2 Whitney disk. This “transfer move” will not create any new unpaired order 1 intersections, and can be iteratively applied to convert all algebraically canceling pairs into geometrically canceling pairs, yielding the desired order 2 framed Whitney tower on A . Description of this modification will be accompanied by Figures 4.7 through Figure 4.11, with some details left to the exercises.

The transfer move starts by pushing p off of W by a finger move into A as shown in Figure 4.8. This finger move is guided by an arc in W along the b -labeled and 1-labeled edges of t_p from p to $\partial_+ W$.

The next step in the transfer move is to push one of the self-intersections of A created by the first finger move across $\partial_+ W'$ by a finger move as in Figure 4.9. This finger move is guided by an arc in A from $\partial_+ W$ to $\partial_+ W'$, and creates a new order 1 intersection $p' \in A \# W'$. (The 4-ball shown in Figures 4.7–4.11 is a neighborhood of the union of this guiding arc with W and W' .) We may ensure that p' has the same sign as p did by choosing this guiding arc to approach from the correct side of $\partial_+ W'$ in A (Exercise 4.8.9), and in fact $t_{p'} = t_p$ by Exercise 4.8.10. As illustrated by the blue Whitney circle in Figure 4.10, p' and q admit an order 2 framed Whitney disk by Exercise 4.8.11.

The construction has also created order 1 intersections $r, s \in A \# A$ which admit a framed embedded Whitney disk V , appearing “underneath” the horizontal sheet in Figure 4.11. This




Figure 4.8




Figure 4.9




Figure 4.10

V has a pair of oppositely-signed intersections with A that admit a framed order 2 Whitney disk whose boundary is indicated in blue in the figure (Exercise 4.8.12).

To see that this transfer move can be iterated to convert all algebraically canceling pairs of intersections into geometrically canceling pairs, observe that disjointly embedded guiding arcs in A can be found for all pairs of Whitney disks, and the new order 1 V -Whitney disks are supported near these arcs and the original Whitney disks' boundaries. The order 2 Whitney disks created in the construction do not need to be controlled since they can only create new intersections of order ≥ 2 .




Figure 4.11

To complete the proof that the vanishing of $\tau_1(A)$ implies that A supports an order 2 framed Whitney tower it just remains to prove the following lemma describing the realization of INT relations for $\mathbb{R}P^2$ s, which was used to achieve all algebraically canceling pairs:

Lemma 4.2. *Let \mathcal{W} be an order 1 framed Whitney tower on A , and let $R : \mathbb{R}P^2 \looparrowright X$ represent $a \in \pi_1 X$ with $a^2 = 1$. Then, after one finger move, A supports an order 1 framed Whitney tower \mathcal{W}' such that $[t(\mathcal{W}')] = [t(\mathcal{W}) + (a, \lambda_0(A, R) + \omega_2(R))] \in \widetilde{\mathcal{T}}_1$.*

Proof. Perform a finger move on A guided by a circle which is isotopic to $R(\mathbb{R}P^1)$ representing a . We may assume that this circle is disjoint from all Whitney disks in \mathcal{W} . The new pair of self-intersections created by this finger move admit a clean local framed Whitney disk W whose boundary is indicated in green in Figure 4.12.




Figure 4.12. After the finger move: A neighborhood of the circle representing a . The thick arrow indicates where the circle extends outside the illustrated local coordinates.

Figure 4.13 shows the preimages in red of a different choice of Whitney arcs that induce the opposite sheet choice as W . (Compare with Figure 4.5, but note that here preimages of the same point are aligned vertically, as opposed to diagonally in Figure 4.5.)

Decomposing $\mathbb{R}P^2$ as the union $Mb \cup_{\partial Mb = \partial D^2} D^2$ of a Möbius band neighborhood Mb around $\mathbb{R}P^1$ and a disk D^2 , we want to use $R(D^2)$ as a Whitney disk W' whose boundary is the image




Figure 4.13

of the red arcs (see Figure 4.14). To do this first observe that we can assume that R restricts to an embedding on Mb . Now extend an isotopy between $R(\mathbb{R}P^1)$ and the finger move circle to a homotopy of $R(\mathbb{R}P^2)$ which restricts to an isotopy of $R(Mb)$. To use the $R(D^2)$ as W' we just need the image of the red arcs to align with $\partial R(Mb)$. This can be accomplished by appropriately choosing the twisting of A in $S^1 \times B^3$ along the finger move circle, which is the core of $R(Mb)$.

By general position we may assume that $A \cap R(\mathbb{R}P^2) \subset R(D^2)$, so $A \natural W' = A \natural R(\mathbb{R}P^2)$. After arranging that $\omega(W') = 0$ by boundary-twists, if needed, it follows that changing \mathcal{W} to \mathcal{W}' by replacing W by W' yields the change $[t(\mathcal{W}')] = [t(\mathcal{W}) + (a, \lambda_0(A, R) + \omega_2(R))] \in \widetilde{\mathcal{T}}_1$. Here the $(a, \omega_2(R))$ term comes from boundary-twisting (and interior-twisting) as needed to make W' framed (as in section 4.2.1, any interior twists contribute trivially in $\widetilde{\mathcal{T}}_1$). \square




Figure 4.14. The image under A of the red arcs in Figure 4.13. Solid arcs are in the present, dotted arcs are in the future, and dashed arcs are in the past.

4.3. Splitting twisted Whitney towers

A framed Whitney tower is *split* if the set of singularities in the interior of any Whitney disk consists of either a single point, or a single boundary arc of a Whitney disk, or is empty. This can always be arranged, as observed in Lemma 13 of [30] (Lemma 3.5 of [26]), by performing finger moves along Whitney disks guided by arcs connecting the Whitney disk boundary arcs (see Figure 1.10). Implicit in this construction is that the finger moves preserve the Whitney disk twistings (by not twisting relative to the Whitney disk that is being split – see Figure 4.15).




Figure 4.15. A neighborhood of a twisted finger move which splits a Whitney disk into two Whitney disks. The vertical black arcs are slices of the new Whitney disks, and the grey arcs are slices of extensions of the Whitney sections. The finger-move is supported in a neighborhood of an arc in the original Whitney disk running from a point in the Whitney disk boundary on the “upper” surface sheet to a point in the Whitney disk boundary on the “lower” surface sheet. (Before the finger-move this guiding arc would have been visible in the middle picture as a vertical black arc-slice of the original Whitney disk.)

A Whitney disk W is *clean* if the interior of W is embedded and disjoint from the rest of the Whitney tower. In the setting of twisted Whitney towers, it simplifies the combinatorics of controlled manipulations to use “twisted” finger moves to similarly split twisted Whitney disks into ± 1 -twisted clean Whitney disks.

We call a twisted Whitney tower *split* if all of its non-trivially twisted Whitney disks are clean and have twisting ± 1 , and all of its framed Whitney disks are split in the usual sense (as for framed Whitney towers).

Recall from section 1.10 our convention of embedding into a Whitney tower the trees associated to unpaired intersections and Whitney disks.

Lemma 4.3 ([8, 30]). *If A supports an order n twisted Whitney tower \mathcal{W} , then A is homotopic (rel ∂) to A' which supports a split order n twisted Whitney tower \mathcal{W}' , such that:*

- (i) *The sub-multiset of signed framed trees $\sum_p \epsilon_p \cdot t_p \subset t(\mathcal{W})$ is isomorphic to the sub-multiset of signed framed trees $\sum_{p'} \epsilon_{p'} \cdot t_{p'} \subset t(\mathcal{W}')$.*
- (ii) *Each $\omega(W_j) \cdot J^\infty$ in $t(\mathcal{W})$ gives rise to exactly $|\omega(W_j)|$ -many $\pm 1 \cdot J^\infty$ in $t(\mathcal{W}')$, where each twisting coefficient ± 1 of the J^∞ in $t(\mathcal{W}')$ has the same sign as the twisting $\omega(W_j)$ of the original $W_j \subset \mathcal{W}$.*

Proof. Illustrated in Figure 4.15 is a local picture of a twisted finger move, which splits one Whitney disk into two, while also changing twistings. If the original Whitney disk in Figure 4.15 was framed, then the two new Whitney disks will have twistings $+1$ and -1 , respectively. In general, if the arc guiding the finger move splits the twisting of the original Whitney disk into ω_1 and ω_2 zeros of the extended Whitney section, then the two new Whitney disks will have twistings $\omega_1 + 1$ and $\omega_2 - 1$, respectively. Thus, by repeatedly splitting off framed corners into ± 1 -twisted Whitney disks, any ω -twisted Whitney disk ($\omega \in \mathbb{Z}$) can be split into $|\omega|$ -many $+1$ -twisted or -1 -twisted clean Whitney disks, together with split framed Whitney disks containing any interior intersections in the original twisted Whitney disk. Combining this with the untwisted splitting [30, Lem.13] of the framed Whitney disks illustrated in Figure 1.10 gives the result. \square

4.4. The Whitney move IHX relation

Suppose that \mathcal{W} is a split Whitney tower on A . Let p be an unpaired intersection with $t_p \subset \mathcal{W}$. Define the *split subtower* $\mathcal{W}_p \subset \mathcal{W}$ to be the union of the Whitney disks containing the trivalent vertices of t_p together with sheets of A around the boundary arcs of each of these Whitney disks whose boundary lies in a sheet of A . These sheets inherit the indices of the order 0 surfaces containing them, so the univalent labels of t_p are unchanged. By construction each Whitney disk in \mathcal{W}_p is framed and embedded.




Figure 4.16. The *framed Whitney move IHX relation* replaces a split subtower whose signed tree looks locally like the one on the left with a pair of nearby disjoint split subtowers whose signed trees look locally like the trees on the right.

Lemma 4.4 (Whitney move IHX relation). *Let \mathcal{W}_p be a split subtower in a split Whitney tower \mathcal{W} , and let $W_{((I,J),K)}$ be a Whitney disk in \mathcal{W}_p so that t_p looks locally like the leftmost tree in Figure 4.16 near $W_{((I,J),K)}$. Then \mathcal{W} can be modified in a regular neighborhood $\nu(\mathcal{W}_p)$ of \mathcal{W}_p yielding a split Whitney tower on the same order 0 surface sheets, with \mathcal{W}_p replaced by disjoint split subtowers $\mathcal{W}_{p'}$ and $\mathcal{W}_{p''}$ contained in $\nu(\mathcal{W}_p)$ such that the signed trees $t_{p'}$ and $t_{p''}$ are as pictured on the right hand side of Figure 4.16.*




Figure 4.17. The Whitney move IHX construction starts with a $W_{(I,J)}$ Whitney move on W_I . Note that the intersection $p = W_{((I,J),K)} \pitchfork W_L$ is *not* shown in this figure (and is suppressed in subsequent figures as well).




Figure 4.18. The intersection point $W_J \cap W'_{(I,K)}$ is ‘transferred’ via a finger move (top) to create a canceling pair $W_J \cap W_{(I,K)}$ paired by $W_{(J,I,K)}$ at the cost of also creating $W_J \cap W_K$ paired by $W_{(J,K)}$ and $W_I \cap W_{(J,K)}$ paired by $W_{(I,J,K)}$ (bottom).

Proof. As a preliminary step, observe that the labelling of the trivalent vertices in the left-hand tree in Figure 4.16 indicates that the unpaired intersection $p \in \mathcal{W}_p$ corresponds to an edge in the L -subtree. By applying the move of Figure 1.18 in section 1.20 we may assume that in fact $p = W_{(I,J,K)} \pitchfork W_L$, so that the edge corresponding to p is the L -labeled edge in the lefthand tree. This will simplify later steps in the construction. For visual clarity p will be suppressed from view in the figures.

Now do the $W_{(I,J)}$ Whitney move on W_I (see Figure 4.17). This eliminates the canceling pair of intersections between W_I and W_J at the cost of creating two canceling pairs of intersections between W_I and W_K which we pair by Whitney disks $W_{(I,K)}$ and $W'_{(I,K)}$ which are meridional disks to W_J as illustrated in Figure 4.17.

The new Whitney disks $W_{(I,K)}$ and $W'_{(I,K)}$ each have a single interior intersection with W_J and the next step is to “transfer” (as illustrated in the upper part of Figure 4.18) the intersection point $W_J \cap W'_{(I,K)}$ to create a canceling pair $W_J \cap W_{(I,K)}$ paired by $W_{(J,I,K)}$ at the cost of also creating $W_J \cap W_K$ paired by $W_{(J,K)}$ and $W_I \cap W_{(J,K)}$ paired by $W_{(I,J,K)}$ (as illustrated in the lower part of Figure 4.18). Note that Figure 4.18 differs from Figure 4.17 by a change of coordinates which brings the sheet of W_K into the “present” slice of 3-space.

This transfer move is combinatorially the same as in section 4.2.2 but here applied to higher-order sheets. The important thing to note here is that the finger move is guided by an




Figure 4.19. The transferring finger move is guided by an arc a (shown in green) which can be taken to run along what used to be the part of the boundary arc of $W_{((I,J),K)}$ lying in W_K . The bottom picture is the same as the top picture but indicating where $W_{((I,J),K)}$ used to be.

arc a (see Figure 4.19) from $\partial W'_{(I,K)}$ to $\partial W_{(I,K)}$ in W_K and we can take this arc to run along what used to be the part of $\partial W_{((I,J),K)}$ lying in W_K . This is illustrated in the lower part of Figure 4.19 which gives a better picture of the situation before the finger move is applied. The Whitney disks $W_{(I,J,K)}$ and $W_{(J,I,K)}$ are taken to be parallel copies of the old $W_{((I,J),K)}$ as follows: The boundary of $W_{(I,J,K)}$ (resp. $W_{(J,I,K)}$) consists of arcs a' and b' (resp. a'' and b''), where a' and a'' are tangential push-offs of a in W_K and b' and b'' are normal push-offs of what was the boundary arc b of $W_{((I,J),K)}$ in $W_{(I,J)}$. This is shown in both Figure 4.20 and Figure 4.21, where again it is easier to picture things before the transferring finger move. Since $W_{((I,J),K)}$ was framed and embedded, $W_{(I,J,K)}$ and $W_{(J,I,K)}$ can be formed from two disjoint parallel copies of $W_{((I,J),K)}$ which each intersect W_L in a single point as $W_{((I,J),K)}$ did.

Thus exactly two new unpaired intersection points p' and p'' have been created (near where p was) with corresponding trees $t(p')$ and $t(p'')$ as shown locally in the right hand side of Figure 4.16. After the transferring finger move, the $W'_{(I,K)}$ Whitney move can be done

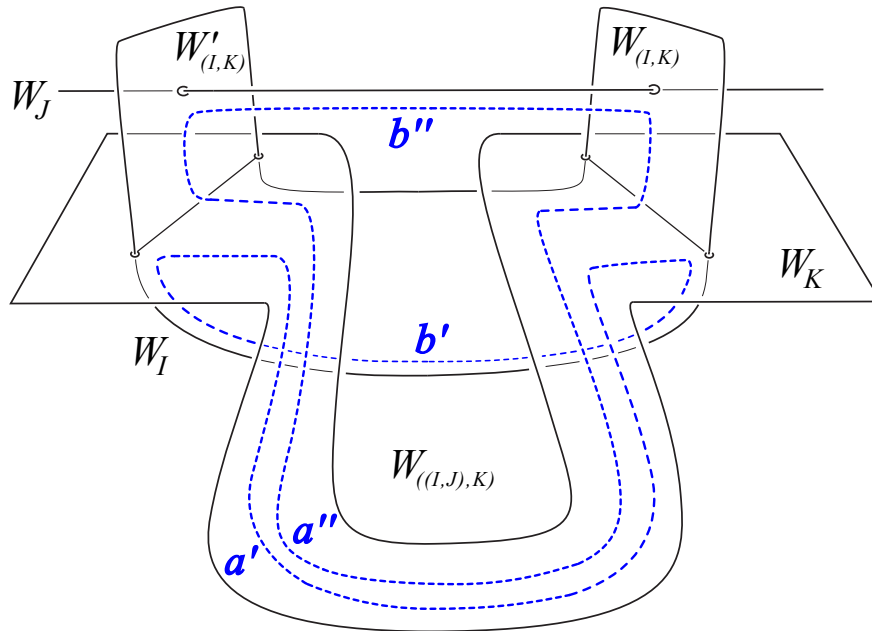


Figure 4.20. Before the transfer move: New Whitney disks $W_{(I,(J,K))}$ and $W_{(J,(I,K))}$, whose boundaries are the unions of arcs $a' \cup b'$ and $a'' \cup b''$ (see also Figure 4.21), will be created from parallel copies of the old $W_{((I,J),K)}$.

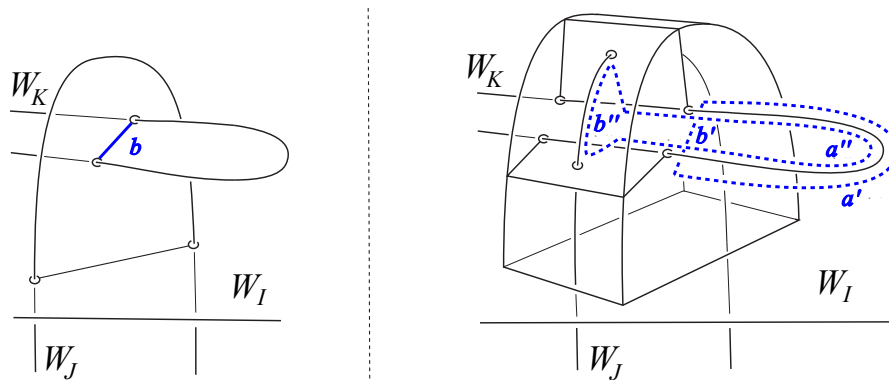


Figure 4.21. Applying the transfer move to the right-hand side will create new Whitney disks $W_{(I,(J,K))}$ and $W_{(J,(I,K))}$, whose boundaries are the unions of arcs $a' \cup b'$ and $a'' \cup b''$ (see also Figure 4.20), from parallel copies of the old $W_{((I,J),K)}$ shown on the left.

(on either sheet) without affecting anything else. Finally, W_I , W_J and W_K will need to be split since they now each contain two boundary arcs of Whitney disks. Splitting W_I , W_J and W_K down into the lower order Whitney disks (as in Figure 1.10 and Lemma 4.3(i)) yields the two split subtowers $\mathcal{W}_{p'}$ and $\mathcal{W}_{p''}$. \square

4.5. The Whitney move twisted IHX relation

Recall the twisted IHX/Jacobi relation in the even order twisted tree groups (section 2.1):

$$I^\infty = H^\infty + X^\infty - \langle H, X \rangle$$

where I , H and X denote rooted trees which correspond to the three terms of a Jacobi identity, i.e. they differ locally as in Figure 4.16, where we now are interpreting the trees in the figure

as being rooted trees with each of the L -subtrees containing a root univalent vertex. Here we are using Roman font in the twisted IHX relation rather than the italic font used in section 2.1 to clarify the distinction between the lefthand term “ I ” of the Jacobi identity and the subtree “ I ” in each of the three terms in Figure 4.16.

We will modify the proof of the above Whitney move IHX relation (Lemma 4.4) to show how to locally replace a twisted Whitney disk whose associated signed ∞ -tree is I^∞ by two twisted Whitney disks whose associated signed ∞ -trees are H^∞ and X^∞ , respectively, together with a single unpaired intersection p having signed tree $-(H, X)$.

Suppose that \mathcal{W} is a split Whitney tower on A . Let W be a Whitney disk in \mathcal{W} with rooted tree $J \subset \mathcal{W}$. Then the *split subtower* $\mathcal{W}_W \subset \mathcal{W}$ is defined to be the union of the Whitney disks containing the trivalent vertices of J together with sheets of A around the boundary arcs of each of these Whitney disks whose boundary lies in a sheet of A . Note that $W \subset \mathcal{W}_W$. If W is twisted, then by construction all the other Whitney disks in \mathcal{W}_W are framed and the twisting of W is ± 1 by the definition of a split Whitney tower, and W_J contains the ∞ -labeled root vertex of $J^\infty \subset \mathcal{W}_W$.

Lemma 4.5. *Let \mathcal{W}_W be a split subtower for an ϵ -twisted Whitney disk W in a split Whitney tower \mathcal{W} , with the ∞ -tree I^∞ associated to W looking locally like the leftmost tree in Figure 4.16 (with the L -subtree containing a root univalent vertex), and denote by H^∞ and X^∞ the ∞ -trees which only differ locally from I^∞ as in the first two trees in the right side of the equation in Figure 4.16.*

Then \mathcal{W} can be modified in a regular neighborhood $\nu(\mathcal{W}_W)$ of \mathcal{W}_W yielding a split Whitney tower on the same order 0 surfaces, with \mathcal{W}_W replaced by disjoint split subtowers $\mathcal{W}_{W'}$, $\mathcal{W}_{W''}$ and \mathcal{W}_p contained in $\nu(\mathcal{W}_W)$ such that:

- (i) W' is an ϵ -twisted Whitney disk with associated ∞ -tree H^∞ , and
- (ii) W'' is an ϵ -twisted Whitney disk with associated ∞ -tree X^∞ , and
- (iii) the unpaired intersection p has sign $-\epsilon$ and associated tree (H, X) .

Proof. We first consider the case where the L -labeled subtree is order zero, which means that L is just the ∞ -label, and the upper trivalent vertex of the I^∞ -tree in Figure 4.16 corresponds to the clean ϵ -twisted $W_{((I,J),K)}$, with ∞ -tree $I^\infty = ((I,J), K)^\infty$. Then the construction in the proof of Lemma 4.4, which starts by performing a Whitney move on the framed Whitney disk $W_{(I,J)}$ corresponding to the lower trivalent vertex of the I -tree, exchanges $W_{((I,J),K)}$ for two Whitney-parallel Whitney disks $W_{(I,(J,K))}$ and $W_{(J,(I,K))}$ (Figures 4.20 and 4.21). In our current setting these two new Whitney disks W' and W'' inherit the ϵ -twisting of W , and have associated twisted trees H^∞ and X^∞ . And because of the twisting $\epsilon = \pm 1$, there is now a single new intersection $p = W' \cap W''$. In the construction of Lemma 4.4 W' inherits the orientation of W , and W'' inherits the opposite orientation, but in our current setting the signs of the twisted trees associated to W and W'' are given by their twisting $\epsilon = \omega(W') = \omega(W'')$ which does not depend on orientations. So we are free to choose orientations on W' and W'' so that the sign of p is $-\epsilon$. Finally, splitting yields the desired split subtowers $\mathcal{W}_{W'}$, $\mathcal{W}_{W''}$ and \mathcal{W}_p .

Now we consider the case where the L -labeled subtree has positive order. This means that the Whitney disk $W_{((I,J),K)}$ is framed and contains a boundary arc of some higher-order Whitney disk V corresponding to a trivalent vertex of L which is adjacent to the trivalent vertex for $W_{((I,J),K)}$ in the I^∞ tree. As above, the construction of Lemma 4.4 exchanges $W_{((I,J),K)}$ for two parallel Whitney disks $W_{(I,(J,K))}$ and $W_{(J,(I,K))}$, but now these three Whitney disks are all framed. The pair of intersections between $W_{((I,J),K)}$ and some sheet W_{L_1} that was paired by V gives rise to two pairs of intersections $W_{L_1} \pitchfork W_{(I,(J,K))}$ and $W_{L_1} \pitchfork W_{(J,(I,K))}$. To continue with the construction we need to find Whitney disks for these new intersection pairs. This will be accomplished using parallels of V , and we state this as a lemma for future reference:

Lemma 4.6. *Let V be an embedded Whitney disk for a pair of intersections between embedded sheets A and B , and let A' and A'' be pairwise disjoint parallel copies of A .*

- (i) *If V is framed, then there exist pairwise disjointly embedded framed Whitney disks V' pairing $A' \pitchfork B$, and V'' pairing $A'' \pitchfork B$, which are parallel to V .*
- (ii) *If V is ϵ -twisted ($\epsilon = \pm 1$), then there exist embedded ϵ -twisted Whitney disks V' pairing $A' \pitchfork B$, and V'' pairing $A'' \pitchfork B$, which are parallel to V , such that $V' \pitchfork V''$ is a single point.*




Figure 4.22. A neighborhood of the framed embedded Whitney disk V pairing sheets A and B contains (in nearby time-slices) parallel Whitney disks V' and V'' pairing the intersections between B and parallels A' and A'' of A . Each of V' and V'' has one boundary arc which is tangent to a boundary arc of V , and the other boundary arc normal to the boundary arc of V in A . Only B extends into all time-slices; the three Whitney disks and other three sheets are each contained in single time slices.

Proof. Figure 4.22 shows the construction of V' and V'' which are Whitney-parallel to V in the case that V is framed. If V is (± 1) -twisted, then Figure 4.22 is only accurate near the Whitney disk boundaries, but since V' and V'' are Whitney parallel it follows that V' and V'' intersect in a single point since they inherit the (± 1) -twisting of V . \square

Returning to the twisted IHX construction, and simplifying notation by writing $A = W_{((I,J),K)}$, and $B = W_{L_1}$, and $A' = W_{(I,(J,K))}$, and $A'' = W_{(J,(I,K))}$, Lemma 4.6 gives us V' pairing $A' \pitchfork B$, and V'' pairing $A'' \pitchfork B$.

If V was the ϵ -twisted W , then we take $W' = V'$, $W'' = V''$, and $p = W' \pitchfork W''$. Splitting then yields the desired split subtowers $\mathcal{W}_{W'}$, $\mathcal{W}_{W''}$ and \mathcal{W}_p .

If V was not the ϵ -twisted W , then V intersected some sheet W_{L_2} in a pair of intersections paired by some Whitney disk V_2 , and hence the parallels V' and V'' also each intersect W_{L_2} in a pair of intersections. Again applying Lemma 4.6 (but now with $A := V_2$ and $B := W_{L_2}$), yields V'_2 and V''_2 pairing $V' \pitchfork W_{L_2}$ and $V'' \pitchfork W_{L_2}$, with V'_2 and V''_2 inheriting the twisting or the intersections that V_2 had.

We continue to apply Lemma 4.6 in this way as needed until reaching the ϵ -twisted $W = V_n$, which gives rise to $W' = V'_n$, $W'' = V''_n$ and $p = W' \pitchfork W''$. Splitting then yields the desired split subtowers $\mathcal{W}_{W'}$, $\mathcal{W}_{W''}$ and \mathcal{W}_p . \square

4.6. Outline of twisted order-raising obstruction theory proof

Recall from sections 2.1, 2.2 and 2.3 the statement of Theorem 2.5: A link $L \subset S^3$ bounds an order n twisted $\mathcal{W} \subset B^4$ with $\tau_n^\infty(\mathcal{W}) = 0 \in \mathcal{T}_n^\infty$ if and only if L bounds an order $n + 1$ twisted Whitney tower.

Theorem 2.5 is essential to the classification of order n twisted Whitney towers in the 4-ball discussed in Section 2, and this section gives a brief outline of the proof, as given in [8,

Thm.1.9]. This proof, which depends on the above Whitney move IHX Lemmas 4.4 and 4.5, generalizes part of the construction of an order 2 framed Whitney tower in section 4.2.

The “if” direction of the theorem holds since by definition any order $n + 1$ twisted Whitney tower is also an order n twisted Whitney tower with no unpaired order n intersections or order $n/2$ twisted Whitney disks. For the “only if” direction, we will sketch how the realization of the relations in \mathcal{T}_n^∞ by geometric constructions can be used to arrange that all unpaired intersections and twisted Whitney disks occur in “algebraically canceling” pairs (representing inverse elements in \mathcal{T}_n^∞), which can then be exchanged for “geometrically canceling” intersection pairs (admitting Whitney disks) and framed Whitney disks of order $n/2$. The strategy is analogous to the order 2 construction in section 4.2, but now we are dealing with twisted Whitney disks and higher-order intersections, and hence more complicated trees (although here in B^4 we do not have to keep track of edge decorations and INT relations).

We can assume that \mathcal{W} is split (Lemma 4.3), and that $t(\mathcal{W})$ contains no framed trees of order $> n$, and no twisted trees of order $> n/2$ (Exercise 2.13.1).

The condition $\tau_n^\infty(\mathcal{W}) = 0 \in \mathcal{T}_n^\infty$ means that in the free abelian group on order n framed trees and order $n/2$ twisted trees $t(\mathcal{W})$ lies in the span of the relators in section 2.1 which define \mathcal{T}_n^∞ . As usual we consider the trees $t(\mathcal{W})$ to be embedded in \mathcal{W} .

As described in [8, Sec.4.1], using also [30, Sec.4] and [7], there are three main steps to constructing an order $n + 1$ twisted Whitney tower from \mathcal{W} :

First, controlled modifications of \mathcal{W} realizing the relators, as discussed in section 2.2, are used to arrange that the order n trees and order $n/2$ ∞ -trees in $t(\mathcal{W})$ all occur in isomorphic oppositely-signed algebraically canceling pairs (see the start of section 4 of [8]).

Secondly, using the above Whitney move IHX Lemmas 4.4 and 4.5, all these paired trees are converted into pairs of “simple” (right- or left-normed) trees by IHX constructions. These simple trees are characterized by the property that every trivalent vertex is adjacent to a univalent vertex. This corresponds to every Whitney disk having a boundary arc on an order zero disk. (The reason for this step will be explained momentarily in the description of the third step.) After this step all the trees still occur in algebraically canceling pairs.

The final third step uses an iterated higher-order variation of the “transfer move” used in section 4.2.2 to achieve geometric cancellation for algebraically canceling order 1 pairs. This step is described in detail for algebraically canceling pairs of framed trees in [30, Lem.15], and for twisted order $n/2$ Whitney disks in [8, Sec.4.1], and requires that all pairs are simple, as arranged in the second step. The reason for this requirement has to do with the connectivity of sheets that is needed for the transfer move. This can be seen by observing that the construction of the two order 2 Whitney disks in Figure 4.11 depends on all three sheets of A being connected. In higher orders it turns out that having just two of these sheets connected suffices to iterate the move finitely many times until eventually terminating with the desired result, provided one starts at an “end” of a simple tree.

After this third step the new layer of order $n + 1$ Whitney disks have uncontrolled intersections, but all of these new intersections are of order $\geq n + 1$. And the construction combining the twisted Whitney disk pairs into framed Whitney disks (Figures 21–22 in [8, Sec.4.1]) creates only new twisted Whitney disks of order $> n/2$, which are supported near the original twisted Whitney disk pairs, along with intersections of order $\geq n$ among these new twisted Whitney disks. Hence an order $n + 1$ twisted Whitney tower has been created.

Analogous order-raising intersection-obstruction theories are described in [8, Sec.4.4] for order n framed Whitney towers, in [30, Thm.6] for non-repeating Whitney towers, and in [12, Thm.6.17] for “ k -repeating” Whitney towers.

4.7. Whitney disk orientations and the AS relation

This section explains why the signs associated to trees for unpaired intersections in a Whitney tower only depend on the orientations of the underlying order 0 surface and the ambient 4-manifold modulo the AS antisymmetry relations.

Let \mathcal{W} be a Whitney tower in an oriented 4-manifold X . The order 0 surface supporting \mathcal{W} comes with a fixed orientation, and arbitrary orientations on the Whitney disks in \mathcal{W} are chosen and fixed. Fix a choice of either the positive or negative corner convention described in section 1.19. Here we will refer to this fixed convention choice as “our corner convention”.

For $p \in W_I \pitchfork W_J$ an unpaired intersection between Whitney disks W_I and W_J in \mathcal{W} , the corresponding tree $t_p = \langle I, J \rangle$ from section 1.10 is embedded in \mathcal{W} in a way that satisfies our corner convention, and the orientations of the trivalent vertices of t_p are taken to be induced by the Whitney disk orientations.

We will consider here the effect on the signed oriented tree $\epsilon_p \cdot t_p$ of switching the fixed orientation choice on any of the Whitney disks corresponding to the trivalent vertices of t_p . The conclusion will be that any such orientation switch corresponds to an AS relation (Figure 2.1), so that modulo AS relations $\epsilon_p \cdot t_p$ only depends on the orientation of the underlying order 0 surface.

First of all, the sign $\epsilon_p = \pm 1$ is determined by comparing the concatenated orientations of W_I and W_J with the orientation of X at p . So in the case that $W_I \neq W_J$, then switching the orientation of one of either W_I or W_J will switch the trivalent orientation of the corresponding trivalent vertex of t_p , and will switch the sign ϵ_p , meaning that $\epsilon_p \cdot t_p$ is changed by an AS relation.

In the case that $W_I = W_J$, then switching the orientation of $W_I = W_J$ will switch the trivalent orientation of t_p at both of the trivalent vertices adjacent to the edge passing through p , and will not change the sign ϵ_p , meaning that $\epsilon_p \cdot t_p$ is changed by two AS relations.

Now consider the effect of switching the orientation on a Whitney disk W_{K_1} other than W_I or W_J , so that W_{K_1} contains a boundary arc of a higher-order Whitney disk $W_{(K_1, K_2)}$ which pairs intersections q and r between W_{K_1} and some W_{K_2} , as in the left side of Figure 4.23.

The orientation switch on W_{K_1} changes the cyclic orientation at the trivalent vertex in W_{K_1} , and it also switches the signs of q and r . This switching of the signs of q and r means that the embedding of t_p needs to be changed near the trivalent vertex v of $W_{(K_1, K_2)}$ in order to preserve our corner convention. The effect of this convention-preserving change in the embedding of t_p near v is to switch the cyclic orientation of t_p at v as shown in the right side of Figure 4.23. Note that the orientation of the Whitney disk $W_{(K_1, K_2)}$ does not change.




Figure 4.23. Before (left) and after (right) the orientation switch on W_{K_1} , using the positive corner convention. Any interior intersection(s) in $W_{(K_1, K_2)}$ are not shown.

Since the orientation of $W_{(K_1, K_2)}$ is not changed, the sign ϵ_p of $p \in W_I \natural W_K$ is unchanged. (It is possible that $W_{(K_1, K_2)} = W_I$ or $W_{(K_1, K_2)} = W_J$.) So the result of switching the orientation of any one W_{K_1} (with $W_{K_1} \neq W_I$, and $W_{K_1} \neq W_J$) while preserving our corner convention is the same as applying *two* AS relations to t_p , one at the trivalent vertex in W_{K_1} , and one at $v \in W_{(K_1, K_2)}$.

4.8. Section 4 Exercises

4.8.1. Exercise

In the proof of independence of $\tau_1(A)$ on Whitney disk interiors given in section 4.1.4 show that the sphere $S = (W \setminus w) \cup (V' \setminus v')$ is homotopic to $S' = W \cup V'$ by a homotopy supported near γ . HINT: Consider a homotopy from S' to S which pulls the common collars $w = v'$ slightly apart while perturbing S' to be smooth, and then shrinks the union of the collars away from γ .

4.8.2. Exercise

Check that the signed tree t_p associated to $p \in \partial_\epsilon W \cap \partial_\delta V$ in Equation (4.1) near Figure 4.1 is well-defined, including the case that $W = V$.

4.8.3. Exercise

Check that in each of the other cases of Figure 4.3 pushing $\partial_\delta W$ into $\partial_\epsilon V$ across the \pm self-intersection of A paired by V creates an algebraically canceling pair of signed trees.

4.8.4. Exercise

In the setting of section 4.1.6 and Figure 4.4, check that the order 1 decorated trees corresponding to the intersections $A \natural V'$ cancel with those corresponding to the oppositely-signed parallel copy of V' in V .

4.8.5. Exercise

In the discussion in section 4.1.7 of the pairing $\lambda_0(A, R)$ for $A : S^2 \natural X$ and $R : \mathbb{RP}^2 \natural X$ with the generator of $\pi_1 \mathbb{RP}^2$ mapping to $a \in \pi_1 X$, show that changing the choice of sheet-changing path through $p \in A \natural R$ changes g_p by right multiplication by a^n and changes ϵ_p by multiplication by $(-1)^n$ for some integer n .

4.8.6. Exercise

From section 4.1.8: For $A : S^2 \natural X$, let p and q be a positive and a negative transverse self-intersection of A , and denote the preimages by $A^{-1}(p) = \{x, x'\} \subset S^2$, and $A^{-1}(q) = \{y, y'\} \subset S^2$.

If p and q have common group element $g_p = a = g_q$, then any Whitney disk W pairing p and q induces a pairing of $\{x, x'\}$ with $\{y, y'\}$ since each arc of ∂W runs between a sheet of A around p and a sheet of A around q .

Check that Whitney disks exist for *both* of the two pairing choices $x \leftrightarrow y, x' \leftrightarrow y'$ and $x \leftrightarrow y', x' \leftrightarrow y$ if and only if $a^2 = 1 \in \pi_1 X$.

4.8.7. Exercise

From section 4.1.8 and Figure 4.5: Check that the union $A(D) \cup W \cup V$ defines the image of a map $R : \mathbb{RP}^2 \rightarrow X$ which sends the generator of $\pi_1 \mathbb{RP}^2$ to $a \in \pi_1 X$.

4.8.8. Exercise

Let W and V be Whitney disks pairing the same intersections p and q between surfaces in a 4-manifold, and suppose that $\partial W = \partial V$ near p and near q . Convince yourself that by an isotopy it can be arranged that $W = V$ near p and near q without creating any new intersections.

4.8.9. Exercise

Check that in the setting of Figures 4.7, 4.8 and 4.9 it can be arranged that p' has the same sign as p .

4.8.10. Exercise

Check that in the setting of Figures 4.7, 4.8 and 4.9 it can be arranged that p' has the same decorated tree as p .

4.8.11. Exercise

Use the previous two exercises to check that p' and q in Figure 4.10 admit an order 2 Whitney disk.

4.8.12. Exercise

Continuing the previous three exercises, the construction of subsection 4.2.2 has also created order 1 intersections $r, s \in A \# A$ which admit a framed embedded Whitney disk V , shown “underneath” the horizontal sheet in Figure 4.11. Check that the pair of transverse intersections between A and the embedded order 1 Whitney disk V in Figure 4.11 admits a framed order 2 Whitney disk (whose boundary is indicated in blue in the figure).

4.8.13. Exercise

In the paragraph before the proof of Lemma 4.1, check that the group element associated to any point in $A(D) \# A(D)'$ is 1 or a .

4.8.14. Exercise

Check that the operation of splitting a Whitney tower preserves edge decorations on trees. Do this first for decorated order 1 trees (as in section 3.4), then generalize to higher-order trees.

4.8.15. Exercise

Check that the signs of the trees created by the IHX construction in the proof of Lemma 4.4 have the correct signs, as shown in the right side of Figure 4.16.

4.8.16. Exercise

Check that the IHX construction in the proof of Lemma 4.4 also works for decorated trees.

4.8.17. Exercise

Prove that $\lambda_1(A_1, A_2, A_3)$ from Theorem 3.2 is a well-defined homotopy invariant which vanishes if and only if $A_1 \cup A_2 \cup A_3$ admits an order 2 non-repeating Whitney tower by adapting

to this easier non-repeating setting the proof given in sections 4.1 and 4.2 of the analogous statements for $\tau_1(A)$.

HINT: First show that $\lambda_1(A_1, A_2, A_3)$ is independent of the choice of order 1 non-repeating Whitney tower, following sections 4.1.4, 4.1.5 and 4.1.6 (but without worrying about Whitney disk twistings in section 4.1.4). Then one gets homotopy invariance exactly as in section 4.1.9. To get an order 2 non-repeating Whitney tower proceed as in section 4.2 to first achieve algebraic cancellation and then geometric cancellation of all order 1 non-repeating intersections.

References

- [1] **R Budney, J Conant, K P Scannell, D Sinha**, *New perspectives on self-linking*. Adv. Math. 191.1 (2005), 78–113.
- [2] **J C Cha**, *Link concordance, homology cobordism, and Hirzebruch-type defects from iterated p -covers*, J. Eur. Math. Soc. 12 (2010) 555–610.
- [3] **J C Cha**, *Rational Whitney tower filtration of links*, Mathematische Annalen (2017) 1–30.
- [4] **J C Cha, K Orr**, *Transfinite Milnor invariants for 3-manifolds*, preprint (2020). arXiv:2002.03208 [math.GT]
- [5] **T Cochran**, *Derivatives of links, Milnor’s concordance invariants and Massey products*, Mem. Amer. Math. Soc. Vol. 84 No. 427 (1990).
- [6] **T Cochran, K Orr, P Teichner**, *Knot concordance, Whitney towers and L^2 -signatures*, Annals of Math., Volume 157 (2003) 433–519.
- [7] **J Conant, R Schneiderman, P Teichner**, *Jacobi identities in low-dimensional topology*, Compositio Mathematica 143 Part 3 (2007) 780–810. arXiv:math/0401427v2 [math.GT]
- [8] **J Conant, R Schneiderman, P Teichner**, *Whitney tower concordance of classical links*, Geom. Topol. (2012) 16 (2012) 1419–1479. arXiv:1202.3463 [math.GT]
- [9] **J Conant, R Schneiderman, P Teichner**, *Tree homology and a conjecture of Levine*, Geom. Topol. 16 (2012) 555–600. arXiv:1012.2780 [math.GT]
- [10] **J Conant, R Schneiderman, P Teichner**, *Universal quadratic forms and Whitney tower intersection invariants*, Proceedings of the Freedman Fest, Geom. Topol. Monographs 18 (2012) 35–60. arXiv:1101.3480 [math.GT]
- [11] **J Conant, R Schneiderman, P Teichner**, *Milnor invariants and twisted Whitney towers*, J. Topology, Volume 7, Issue 1, (2013) 187–224. arXiv:1102.0758 [math.GT]
- [12] **J Conant, R Schneiderman, P Teichner**, *Clasper concordance, Whitney towers and repeating Milnor invariants*, preprint (2020). arXiv:2005.05381 [math.GT]
- [13] **M Freedman, F Quinn**, *The topology of 4-manifolds*, Princeton Math. Series 39 Princeton, NJ, (1990).
- [14] **M Freedman, P Teichner**, *4-manifold topology II: Dwyer’s filtration and surgery kernels.*, Invent. Math. 122 (1995) 531–557.
- [15] **T G Goodwillie, M Weiss**, *Embeddings from the point of view of immersion theory II*, Geom. Topol. 3 (1999) 103–118. doi: 10.2140/gt.1999.3.103
- [16] **N Habegger, G Masbaum**, *The Kontsevich integral and Milnor’s invariants*, Topology 39 (2000) 1253–1289.
- [17] **M Kervaire, J Milnor**, *On 2-spheres in 4-manifolds*, Proc. Nat. Acad. Sci. Vol. 47 (1961) 1651–1657.
- [18] **S Konyagin, M Nathanson**, *Sums of products of congruence classes and of arithmetic progressions*, Int. J. Number Theory Vol. 5 Issue 4 (2009) 625–634.
- [19] **D Kosanović**, *Embedding Calculus and Grope Cobordism of Knots*, preprint (2020) arXiv:2010.05120 [math.GT]
- [20] **Y Matsumoto**, *Secondary intersectional properties of 4-manifolds and Whitney’s trick*, Proceedings of Symposia in Pure mathematics Vol. 32 Part 2 (1978) 99–107.
- [21] **J-B Meilhan, A Yasuhara**, *Characterization of finite type string link invariants of degree < 5* , Math. Proc. Cambridge Phil. Soc. 148:03, 439 (2010).
- [22] **J Milnor**, *Link groups*, Annals of Math. 59 (1954) 177–195.
- [23] **J Milnor**, *Isotopy of links*, Algebraic geometry and topology, Princeton Univ. Press (1957).
- [24] **K Orr**, *Homotopy invariants of links*, Invent. Math. 95 17 (1989) 379–394.
- [25] **M Powell, A Ray, P Teichner**, *The 4-dimensional disc embedding theorem and dual spheres*, preprint (2020) arXiv:2006.05209v1 [math.GT]
- [26] **R Schneiderman**, *Whitney towers and Gropes in 4-manifolds*, Trans. Amer. Math. Soc. 358 (2006) 4251–4278. arXiv:math/0310303 [math.GT]
- [27] **R Schneiderman**, *Simple Whitney towers, half-gropes and the Arf invariant of a knot*, Pacific J. Math. Vol. 222 No. 1, Nov (2005) 169–184. arXiv:math/0310304 [math.GT]
- [28] **R Schneiderman**, *Stable concordance of knots in 3-manifolds*, Alg. and Geom. Topology 10 (2010) 37–432. arXiv:0812.4696 [math.GT]

- [29] **R Schneiderman, P Teichner**, *Higher order intersection numbers of 2-spheres in 4-manifolds*, Alg. and Geom. Topology 1 (2001) 1–29. arXiv:math/0008048 [math.GT]
- [30] **R Schneiderman, P Teichner**, *Whitney towers and the Kontsevich integral*, Proceedings of a conference in honor of Andrew Casson, UT Austin 2003, Geometry and Topology Monograph Series, Vol. 7 (2004) 101–134. arXiv:math/0401441 [math.GT]
- [31] **R Schneiderman, P Teichner**, *Pulling Apart 2-spheres in 4-manifolds*, Documenta Math. 19 (2014) 941–992. arXiv:1210.5534 [math.GT]
- [32] **R Schneiderman, P Teichner**, *Homotopy versus isotopy: spheres with duals in 4-manifolds*, preprint (2019) arXiv:1904.12350 [math.GT]
- [33] **R Stong**, *Existence of π_1 -negligible embeddings in 4-manifolds: A correction to Theorem 10.5 of Freedman and Quinn*, Proc. of the A.M.S. 120 (4) (1994) 1309–1314.
- [34] **M Weiss**, *Embeddings from the point of view of immersion theory I*, Geom. Topol. 3 (1999), 67–101. doi: 10.2140/gt.1999.3.67
- [35] **M Yamasaki**, *Whitney’s trick for three 2-dimensional homology classes of 4-manifolds*, Proc. Amer. Math. Soc. 75 (1979) 365–371.

Dept. of Mathematics, Lehman College, City University of New York, Bronx, NY, USA •
robert.schneiderman@lehman.cuny.edu