Whitney towers and the Kontsevich integral

Rob Schneiderman and Peter Teichner

Courant Institute of Mathematical Sciences, New York University, 251 Mercer Street New York NY 10012-1185

Department of Mathematics, University of California, Berkeley, CA 94720-3840

Email: `schneiderman@courant.nyu.edu` and `teichner@math.berkeley.edu`

Abstract We continue to develop an obstruction theory for embedding 2–spheres into 4–manifolds in terms of Whitney towers. The proposed intersection invariants take values in certain graded abelian groups generated by labelled trivalent trees, and with relations well known from the 3–dimensional theory of finite type invariants. Surprisingly, the same exact relations arise in 4 dimensions, for example the Jacobi (or IHX) relation comes in our context from the freedom of choosing Whitney arcs. We use the finite type theory to show that our invariants agree with the (leading term of the tree part of the) Kontsevich integral in the case where the 4–manifold is obtained from the 4–ball by attaching handles along a link in the 3–sphere.

AMS Classification 57M99; 57M25

Keywords 2–sphere, 4–manifold, link concordance, Kontsevich integral, Milnor invariants, Whitney tower

Dedicated to Andrew Casson on the occasion of his 60th birthday

1 Introduction

Two of Andrew Casson's wonderful contributions to topology were his work on *flexible handles* (now called *Casson towers*) in 4-manifolds, and his invariant for homology 3–spheres, counting representations into $SU(2)$. In this paper we will describe an obstruction theory for disjointly embedding collections of 2–spheres (or 2–disks with fixed boundary) into a 4–manifold that provides a connection between these two aspects of Casson's work. This connection is somewhat indirect, otherwise our paper would be called *Casson towers and the Casson invariant*. In other words, we shall switch from Casson towers to Whitney towers, and from the Casson invariant to the Kontsevich integral. It would be very satisfying to find a more straightforward relationship between Casson's two contributions.

To explain the connection, recall that the Casson invariant is the lowest order (nontrivial) *finite type invariant* of homology 3–spheres. These finite type invariants take values in certain graded abelian groups generated by trivalent graphs. Being *invariants*, they measure the *uniqueness* of 3–manifolds or links in 3–manifolds. We shall explain how similar graphs, better, unitrivalent trees, arise in *existence* questions for 4–manifolds or surfaces in 4–manifolds. It is not totally surprising that raising the dimension by one takes uniqueness to existence questions, after all an isotopy of, say, a knot in a 3–manifold M^3 is nothing but a certain annulus in the 4–manifold $M \times I$. However, the details of such a translation from one dimension to the next are not at all obvious.

In the easiest setting one would like to find obstructions for making the images of maps $A_i : (D^2, S^1) \to (X^4, \partial X)$ *disjoint*, without changing the homotopy classes (and without trying to embed the A_i). In fact, Casson's main Embedding Theorem in [2, Lecture 1] is an example of a special case of this problem: Casson showed that if X is simply connected, all intersection numbers between the A_i vanish, and A_i have *algebraic dual spheres*, then the problem has a positive solution. He used inverses of the Whitney move, now known as Casson or finger moves, to introduce many self-intersections, while trivializing the fundamental group of the complement of one disk at a time (and hence enabling the other disks to be mapped *disjointly*). He then went on to construct Casson towers (with prescribed boundary circles) by iterating the procedure indefinitely, using the fact that the complement of a finite height Casson tower can be made simply connected. These ideas inspired Mike Freedman who proved in [10] that a neighborhood of a Casson tower actually contains an embedded flat disk.

The presence of algebraic dual spheres in Casson's theorem comes from the fact that the proposed application was to the s-cobordism theorem and to the exactness of the surgery sequence in dimension 4. Indeed, Freedman's theorem implies these results in the topological category (for *good* fundamental groups).

There is a more general context in which disjoint maps of disks or spheres can be constructed, namely in the presence of a *non-repeated Whitney tower* (of sufficiently high order), see Theorem 3 below and [31]. The first order stage of this Whitney tower is guaranteed by the vanishing of the intersection numbers whereas the existence of the higher order stages are obstructed by our new proposed invariants. They take values in certain graded abelian groups generated by trivalent trees, which are basically the spines of the Whitney towers. The difference between a Casson tower and a Whitney tower is that in the latter, fewer disks are attached at each stage: In a Casson tower, every intersection point p leads to a new disk (with boundary an arc leaving on one sheet at p and arriving at the other sheet), whereas a Whitney tower only has a new disk for certain *pairs* of intersection points. In particular, it is usually only possibly to find Casson towers in simply connected 4–manifolds, whereas Whitney towers are not restricted by the fundamental group. In fact, in our theory the fundamental group leads to a decoration of the trivalent trees in question, thus giving a much bigger variety of possible obstructions. In addition, Freedman's reimbedding theorem shows that a Casson tower of height 3 already contains an embedded flat disk. However, there are Whitney towers of arbitrary order *not* containing disks, which explains the use of these "weaker" towers in an obstruction theory.

Our Theorem 3 implies Casson's result because algebraic dual spheres can be used to construct non-repeating Whitney towers of arbitrary order. This is already implicit in [11], so our main contribution is a theory *in the absence* of algebraic dual spheres. For example, this applies to concordance questions for links in 3–space. In this context we prove in Theorem 4 below that our invariants agree rationally with (the leading term of) the tree part of the Kontsevich integral, which is the universal finite type concordance invariant [17]. This relates our obstruction theory to the finite type theory and, in particular, to the Casson invariant. It should be mentioned here that Habegger and Masbaum show in [17] that (the leading term of) the tree part of the Kontsevich integral carries exactly the same information as Milnor's $\overline{\mu}$-invariants which were first observed to be concordance invariants by Casson in [3]. Reversing the logic, we have found a 4–dimensional geometric interpretation of this part of the Kontsevich integral, in terms of higher order intersections among Whitney disks. See [8] for an interpretation in terms of *gropes* in 3–dimensions which is stronger in the sense that it works for (the leading term of) the Kontsevich integral, not just of the tree part.

At the time of writing, the setting of Theorem 4 is actually the only case where we have a proof that our intersection invariant is independent of the choice of a Whitney tower, but see Conjecture 1. What we do prove in Theorem 2 is that the vanishing of our intersection invariant for a Whitney tower of order n enables one to build a Whitney tower of the next order $(n+1)$. In that sense, we are producing an obstruction theory since disjointly embedded sheets A_i allow Whitney towers of arbitrary order.

We close this introduction by pointing out that the Whitney towers used in this paper are generalizations of the ones in [5] in that disks of higher order are here allowed to intersect previous stages, as long as these intersection points are paired up by Whitney disks (up to the desired order). In our language, the distinction is made in terms of saying that these Whitney towers have an *order* whereas the Whitney towers of [5] (where different order Whitney disks don't intersect) have a *height*. This is the precise analogue of *class* versus *height* in the theory of gropes, see e.g. [33], ultimately coming from the distinction between the lower central series and the derived series of a group. The latter explains why Whitney towers with a height carry more subtle information. In fact, they are *not* related to the usual finite type theory and hence it is much more difficult to define an obstruction theory. At present, such a theory only exists for knot concordance [5], [6] (using von Neumann signatures to prove nontriviality) and it would be extremely interesting to develop it more generally, i.e. in the context of 2–spheres in 4–manifolds.

2 Statement of Results

We continue to develop the obstruction theory for embedding 2–spheres into 4–manifolds started in [30]. To fix notation, let X be a 4–manifold and $A_1, \dots, A_m$ be generic immersions of 2–spheres (or 2–disks with fixed boundary) into X. We shall work in the smooth setting, even though the techniques of [11] allow a generalization of our work to locally flat surfaces in a topological manifold. The goal is to construct obstructions for changing the A_i, in their regular homotopy class, to embeddings with disjoint images. This is already a very interesting problem for $m = 1$ but we shall not restrict to this case.

The first, well known, invariants are the Wall intersection "numbers" [34]

$$\lambda(A_i, A_j) = \sum_{p \in A_i \cap A_j} \epsilon_p \cdot g_p \quad \in \mathbb{Z}\pi, \quad \pi := \pi_1 X.$$

These count how often A_i and A_j intersect algebraically, including a group element $g_p \in \pi$ and a sign ϵ_p for each intersection point. Similarly, there are self-intersection numbers $\mu(A_i)$ which are well defined only in a certain quotient of the group ring, see below. Recall that in higher dimensions (where A_i are k-spheres, $k > 2$ and X is $2k$-dimensional) the vanishing of these invariants implies that after a finite sequence of Whitney moves [35] the A_i can be represented by disjoint embeddings. In dimension 4, there are well known problems to this procedure (since $2+k = 2k$ for $k = 2$), the most important one being that, generically, the Whitney disks intersect the 2–spheres A_i. The first precise statement concerning the failure of the Whitney trick in dimension 4 was given by Kervaire and Milnor in [18].

In [30] we assumed that these primary intersection numbers vanish which means geometrically that all intersections and self-intersections can be paired by Whitney disks: For each pair of intersection points between A_i and A_j (if $i = j$ these are self-intersections), chose one *Whitney arc* on A_i and one on A_j connecting these two points. Since the fundamental group is controlled in Wall's invariant, the two Whitney arcs together form a null homotopic circle in the ambient 4–manifold, which hence bounds a disk, the *Whitney disk*. Using a choice for such disks, one for each pair of intersection points, we constructed a secondary invariant

$$\tau(A_i, A_j, A_k) \in \mathbb{Z}\pi \times \mathbb{Z}\pi / \dots$$

which measures how the Whitney disks intersect the spheres A_i. Here the indices i, j, k may be repeated, obtaining several slightly distinct geometrical cases just like for Wall's invariants. We recall that by standard procedures the Whitney disks can always be assumed to be disjointly embedded (and framed), and that the only thing which hinders a successful Whitney move is the fact that they are in general not disjoint from the original spheres A_i.

We will first explain a way to unify the above invariants, then suggest a vast generalization and finally discuss a relation to Milnor invariants and the Kontsevich integral (for classical links). For this purpose, assume that the A_i intersect and self-intersect generically, and call the collection $A_1, \dots, A_m$ a *Whitney tower of order 0*. Similarly, if Wall's invariants vanish, and one has chosen generic Whitney disks W_I which pair all intersections and self-intersections of the A_i then one obtains a Whitney tower of order 1. If the τ-invariants vanish, then one can chose Whitney disks for all the intersections of the A_i with the W_I to obtain a Whitney tower of order 2. This procedure can be continued and we give a precise definition of a *Whitney tower of order* n in Section 3. This definition includes orientations of all the surfaces $A_i, W_I, \dots$ in the tower, as well as base points on these surfaces together with whiskers connecting these base points to the base point of X.

2.1 The intersection tree $\tau_n(\mathcal{W})$

Our first observation is that one can canonically associate to each *unpaired* intersection point p of a Whitney tower $\mathcal{W}$ a *decorated* unitrivalent tree t_p of *order* n. The order is the number of trivalent vertices and the decoration is as follows: the univalent vertices of t_p are labelled by the A_i or more abstractly, by $i \in \{1, \dots, m\}$, the edges are labelled by elements from the fundamental group π, and the edges and trivalent vertices are oriented. The tree t_p sits naturally as a subset of $\mathcal{W}$ (Figure 1, details in Section 3) with each trivalent vertex lying in a Whitney disk and each univalent vertex lying in some A_i. Each edge of t_p is a sheet-changing path between vertices in adjacent surfaces, with the

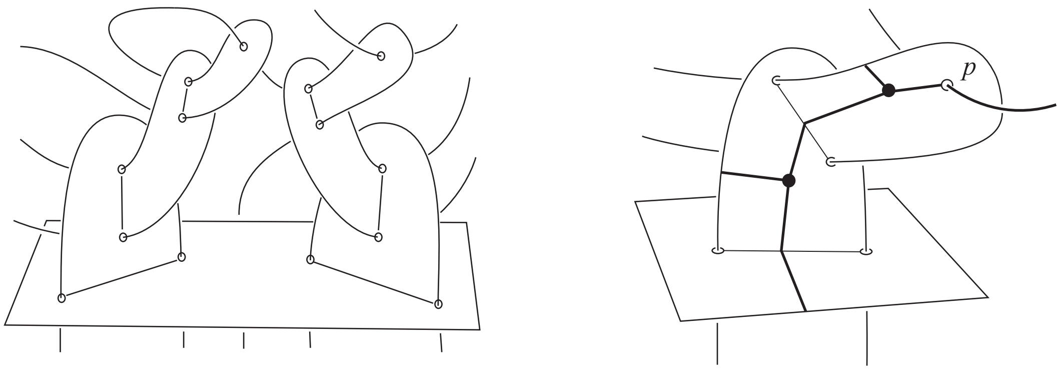

Figure 1: Part of a Whitney tower (left), and part of the unitrivalent tree t_p associated to an unpaired intersection point p in a Whitney tower (right).

group element labelling the edge determined by the loop formed from the path together with the whiskers on the adjacent surfaces. For example, in a Whitney tower of order 0, any intersection point p between A_i and A_j has order 0 and gives a tree t_p consisting of a single edge whose univalent vertices (labelled by i and j) correspond to basepoints in A_i and A_j. This edge is labelled by the group element g_p determined by a loop formed from the whiskers on A_i and A_j together with a path that changes sheets at p where the orientation of the edge corresponds to the direction of the path. For intersection points of order 1 in an order 1 Whitney tower, one gets decorated Y-trees with one trivalent vertex and three univalent vertices labelled by i, j, k (which can repeat).

The central point of this paper is that in an order n Whitney tower $\mathcal{W}$ the trees that correspond to the (unpaired) order n intersection points of $\mathcal{W}$ represent a "higher order" obstruction to homotoping (rel boundary) the A_i to disjoint embeddings. Just like the intersection number $\lambda(A_i, A_j)$ is a sum over all intersection points between A_i and A_j, we define the *intersection tree* $\tau_n(\mathcal{W})$ of an order n Whitney tower $\mathcal{W}$ to be

$$\tau_n(\mathcal{W}) := \sum_p \epsilon_p \cdot t_p \quad \in \mathcal{T}_n(\pi, m).$$

The sum is taken over all order n intersection points p in $\mathcal{W}$ and we consider this sum as taking values in the free abelian group generated by (isomorphism classes of) decorated trees as above, modulo several relations that are motivated geometrically (explained briefly below and in detail in Section 3, particularly

Section 3.8). We denote this quotient by

$$\mathcal{T}(\pi, m) = \bigoplus_{n=0}^{\infty} \mathcal{T}_n(\pi, m),$$

where the order n is the number of trivalent vertices and the univalent labels come from $\{1, \ldots, m\}$, possibly repeated. If this index set is undetermined (or unimportant) we shall just write $\mathcal{T}_n(\pi)$.

The order 0 trees are just single edges and it turns out that

$$\mathcal{T}_0(\pi, 1) \cong \mathbb{Z}\pi/\langle \bar{g} - g \rangle, \quad \bar{g} := w_1(g) \cdot g^{-1},$$

where $w_1 : \pi \to \mathbb{Z}/2$ is the first Stiefel-Whitney class of the ambient 4-manifold. The quotient comes from the fact that an edge with two identical labels has an additional symmetry which changes the orientation of the edge. Moreover, our invariant τ_0 gives exactly Wall's self-intersection invariant μ. To get Wall's intersection number $\lambda(A_1, A_2)$ we just need to evaluate τ_0 in order 0 with *exactly* two labels $1, 2$. The invariants τ from [30] are exactly $\tau_1(\mathcal{W})$ in the various versions of $\mathcal{T}_1(\pi)$, depending on the allowed labels.

A short discussion of the relations in $\mathcal{T}(\pi)$ is in order. They reflect the various choices made in the construction of the Whitney tower, as will be discussed in Section 3 (see also Figure 7 in Section 3). As a consequence, working *modulo* these relations makes our intersection tree τ_n *independent* of the choices below.

- Changing orientations on Whitney disks gives AS, *antisymmetry* relations; they introduce a sign when the cyclic ordering of a trivalent vertex is switched.
- Changing the orientation of an edge changes the label g to $\bar{g}$, the OR *orientation* relation.
- Changing the whiskers gives HOL, *holonomy* relations; they multiply the labels of 3 edges coming into a trivalent vertex by a group element.
- Changing the choice of Whitney arcs, i.e. of the boundaries of Whitney disks, gives the IHX relations.

The last type of relations, well known in dimension 3, is maybe the most surprising aspect of our 4–dimensional theory. We feel that our explanation in terms of the indeterminacy of Whitney arcs is very satisfying [9]. It should be pointed out that graded abelian groups like $\mathcal{T}(\pi)$ arose independently in the 3–dimensional work of Garoufalidis, Kricker and Levine [14], [15]. They study trivalent graphs (instead of unitrivalent trees) and π is usually a 3–manifold group. In some form, the Kontsevich integral gives invariants of links (or 3–manifolds) with values in such graded abelian groups. So these are invariants for the *uniqueness* of 3–dimensional objects, whereas our invariants measure *existence* of 4–dimensional things. In that sense, it might not come as a surprise that there is an overlap between these theories. Note that the restriction to trees is a well known feature if one wants *concordance invariants* in the 3–dimensional context, see [8] or [17].

To make it possible that the intersection tree $\tau_n(\mathcal{W})$ only depends on the A_i, it is in fact necessary to introduce two more types of relations which correspond to changing the choices of Whitney disks (for fixed choices of boundaries:

- The INT *interior or intersection* relations come from the choice of the *interiors* of Whitney disks (which can be changed by summing into any 2–spheres). More generally, they measure indeterminacies coming from certain lower order intersection trees for Whitney towers on subsets of the A_i together with other 2–spheres. A special case of these relations will be examined in detail in [31].
- The FR *framing* relations are generated by certain 2–torsion elements which correspond to manipulations of the interiors of Whitney disks that affect their normal framings. This will be described in [32] but see Figure 2.

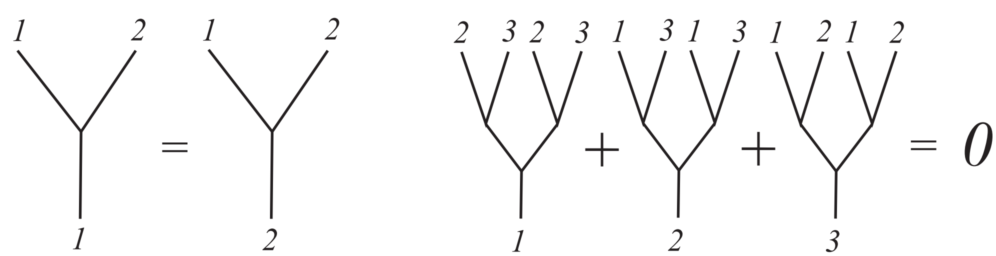

Figure 2: FR relations in order one and three (in a simply-connected 4–manifold).

The INT relations are more subtle in that they actually depend on the ambient 4–manifold X, rather than just on its fundamental group. Both, INT and FR relations will not play a role in this paper, however we will provide evidence supporting the following conjecture by proving a closely related special case.

Conjecture 1 *The intersection tree* $\tau_n(\mathcal{W}) \in \mathcal{T}_n(\pi, m)/\mathrm{INT}, \mathrm{FR}$ *is independent of the choice of the Whitney tower* $\mathcal{W}$*. In fact, it only depends on the regular homotopy classes of the original maps* A_i*, and should be written as* $\tau_n(A_1, \ldots, A_m)$*.*

This result is well known in the Wall case, i.e. for $n = 0$, and it was proven in general for $n = 1$ in [30] (and previously in the simply connected case for $n = 1$ in [26] and [11]).

The following result reflects the obstruction theoretic nature of the intersection tree τ.

Theorem 2 *Let A_i be properly immersed simply-connected surfaces in a 4–manifold, or connected surfaces in a simply-connected 4–manifold. If $\mathcal{W}$ is an order n Whitney tower on the A_i with vanishing intersection tree $\tau_n(\mathcal{W}) \in \mathcal{T}_n(\pi, m)$, then there is an order $(n+1)$ Whitney tower on maps A_i' which are regularly homotopic (rel boundary) to A_i.*

Theorem 2 will be proved in Section 4.

2.2 Immersions with disjoint images

A special case of our invariant only counts those trees t_p whose univalent labels are non-repeating, which means that the number m of spheres A_i is two more than the order n of the intersection point p, $m = n + 2$. Geometrically, one wants to totally ignore *self*-intersections of the spheres A_i and in fact none of the (higher order analogues of) self-intersections in the Whitney tower are paired up. This leads to the notion of a *non-repeated Whitney tower* $\mathcal{W}$ which has also a *non-repeated intersection tree* $\lambda(\mathcal{W})$ that generalizes the λ–invariant of Wall's intersection form. We shall explain these notions in a different paper [31] where we also prove the following beautiful application of the theory.

Theorem 3 *If the 2-spheres $A_1, \ldots, A_{n+2}$ admit a non-repeated Whitney tower $\mathcal{W}$ of order n, such that $\lambda(\mathcal{W})$ vanishes in $\mathcal{T}_n(\pi, n+2)$, then the homotopy classes (rel boundary) of the A_i can be represented by immersions with disjoint images.*

Again, this result was well known for $n = 0$ (see e.g. [20]), and was proven for $n = 1$ in [30] (and for trivial fundamental group in [36]). In the special case discussed in the next section, this result says that a link L in S^3 has vanishing non-repeating Milnor invariants if and only if it bounds disjoint immersions of disks in D^4. In fact, this singular concordance can then be improved to a *link homotopy* from L to the unlink ([13], [12]). This is Milnor's original theorem [24].

2.3 Relation to Milnor invariants and the Kontsevich integral

Given a link L in S^3, there are unique homotopy classes (rel boundary) $A_i : D^2 \to D^4$ of immersions extending L. Therefore, the previous discussion should apply to give link invariants via Whitney towers. The *reduced* Kontsevich integral $Z^t(L)$ is the tree part of the Kontsevich integral of L and in [17] Habegger and Masbaum have shown that the *first non-vanishing term* of $Z^t(L) - 1$ carries exactly the same information as the first non-vanishing Milnor invariants $\mu(L)$. These are the Milnor invariants with repeating indices, also denoted $\bar{\mu}$-invariants [25]. We shall not make this distinction and we consider only the "first non-vanishing" invariants. In the general case one needs to consider *string* links [17].

Denote by $K_n(L)$ the order n term of $Z^t(L) - 1$. Now observe that $K_n(L)$ takes values exactly in $\mathcal{T}_n(m) \otimes \mathbb{Q}$, where m is the number of components of L and the *order* n is the number of trivalent vertices. Here the relations in $\mathcal{T}(m)$ simplify dramatically because $\pi_1(D^4) = 0 = \pi_2(D^4)$ and in fact they reduce to exactly the AS and IHX relations used in the usual definition of the Kontsevich integral. We note that the most commonly used degree in papers on the Kontsevich integral is one half the total number of vertices. For unitrivalent trees, this degree is one more than the number of trivalent vertices, i.e. one *more* than the order that we are using here.

For an oriented link $L \subset S^3$, consider the following four statements.

(i) L bounds a Whitney tower of order n in D^4.

(ii) L bounds disjointly embedded framed gropes of class $(n+1)$ in D^4.

(iii) L has vanishing μ-invariants of length $\leq (n+1)$.

(iv) All terms in $Z^t(L) - 1$ having order $\leq (n-1)$ vanish.

Then (i) is equivalent to (ii) by [28], (iii) is equivalent to (iv) by [17], and (ii) implies (iii) by [23].

The following theorem gives the relation between the Kontsevich integral and our intersection tree τ in the context of the above results.

Theorem 4 *If L bounds a Whitney tower $\mathcal{W}$ of order n in D^4, then*

$$K_n(L) = \tau_n(\mathcal{W}) \quad \in \mathcal{T}_n(m) \otimes \mathbb{Q}$$

which shows that rationally, $\tau_n(L) := \tau_n(\mathcal{W})$ only depends on (the concordance class of) L and can be used to calculate the first non-vanishing terms of the reduced Kontsevich integral as well as the Milnor invariants.

Remark 5 In [32] we shall explain a direct geometric relation between our intersection trees and Milnor's invariants, completely avoiding the Kontsevich integral.

Remark 6 In the nonrepeating case, the groups $\mathcal{T}_n(n+2)$ are torsionfree, and hence tensoring with $\mathbb{Q}$ does not lose any information. This implies our above Conjecture 1 for this very special case (since the FR and INT relations are trivial). By results in [20], Theorem 4 also implies the conjecture for the 2–spheres in the simply connected 4–manifold formed by attaching 0-framed 2–handles to the 4–ball along L in the nonrepeating case (or rationally in the repeating case). It is not unreasonable to believe that the groups $\mathcal{T}_{2n}(m)$ are also torsionfree (with repeated labels allowed). Note that $\mathcal{T}_1(1) \cong \mathbb{Z}/2$ which corresponds exactly to the Arf invariant of a knot (see [26], [29], [30]) and hence shows that statement (iv) does *not* imply (i) in the above theorem. In general, the FR relations are non-trivial for odd orders as will be explained in [32]; see Figure 2 for an example.

3 Whitney towers and intersection trees

The goal of this section is to define the n-th order intersection tree $\tau_n(\mathcal{W})$ of an order n Whitney tower $\mathcal{W}$ in an oriented 4–manifold X. After giving the precise definition of a Whitney tower $\mathcal{W}$, an indexing of the surfaces in $\mathcal{W}$ is given in terms of bracketings and rooted trees which are labelled, oriented and then decorated by elements of the fundamental group $\pi := \pi_1 X$. The unrooted decorated tree t_p associated to an intersection point p in $\mathcal{W}$ then corresponds to a pairing of the rooted trees associated to the intersecting surfaces. Finally, $\tau_n(\mathcal{W})$ is defined as a signed sum of the t_p in the group $\mathcal{T}_n(\pi, m)$, see Section 3.8.

3.1 Whitney towers

We assume our 4–manifolds are oriented and equipped with a basepoint. The reader is referred to [11] for details on immersed surfaces in 4–manifolds, including *Whitney moves* and (Casson) *finger moves*. For more on Whitney towers see [9], [28], [29].

Definition 7

- A *surface of order 0* in a 4–manifold X is a properly immersed surface (boundary embedded in the boundary of X and interior immersed in the interior of X). A *Whitney tower of order 0* in X is a collection of order 0 surfaces.
- The *order of a (transverse) intersection point* between a surface of order n_1 and a surface of order n_2 is $n_1 + n_2$.
- The *order of a Whitney disk* is $n+1$ if it pairs intersection points of order n.
- For $n \geq 0$, a *Whitney tower* $\mathcal{W}$ *of order* $n+1$ is a Whitney tower of order n together with Whitney disks pairing all order n intersection points of $\mathcal{W}$. These top order disks are allowed to intersect each other as well as lower order surfaces.

The Whitney disks in a Whitney tower are required to be *framed* ([11]) and have disjointly embedded boundaries. Intersections in surface interiors are assumed to be transverse. A Whitney tower is *oriented* if all its surfaces (order 0 surfaces *and* Whitney disks) are oriented. A *based* Whitney tower includes a chosen basepoint on each surface (including Whitney disks) together with a *whisker* (arc) for each surface connecting the chosen basepoints to the basepoint of the ambient 4–manifold.

Some further terminology: If $\mathcal{W}$ is an order n Whitney tower containing A_i as its order 0 surfaces then the A_i are said to *admit* an order n Whitney tower and we say that $\mathcal{W}$ is a Whitney tower *on* the A_i.

3.2 Rooted trees and brackets

Non-associative ordered *bracketings* of elements from some index set correspond to *rooted labelled vertex-oriented unitrivalent trees* as follows. Here *rooted* means "having a preferred univalent vertex" (the *root*), *labelled* means that each non-root univalent vertex is labelled by an element from the index set and *vertex-oriented* means that each trivalent vertex is equipped with a cyclic ordering of its incident edges. The *order* of a tree is the number of trivalent vertices.

A bracketing (i) of a singleton element i from the index set corresponds to the rooted order 0 tree $t(i)$ consisting of a single edge with one vertex labelled by i and the other vertex designated as the root. A bracketing (I, J) of brackets I and J corresponds to the *rooted product* $t(I, J) := t(I) * t(J)$ of the trees $t(I)$

and $t(J)$ which identifies together the roots of $t(I)$ and $t(J)$ to a single vertex and "sprouts" a new rooted edge at this vertex (Figure 3) with the cyclic order at the new trivalent vertex given by taking the edges coming from I, J and the root in that order.

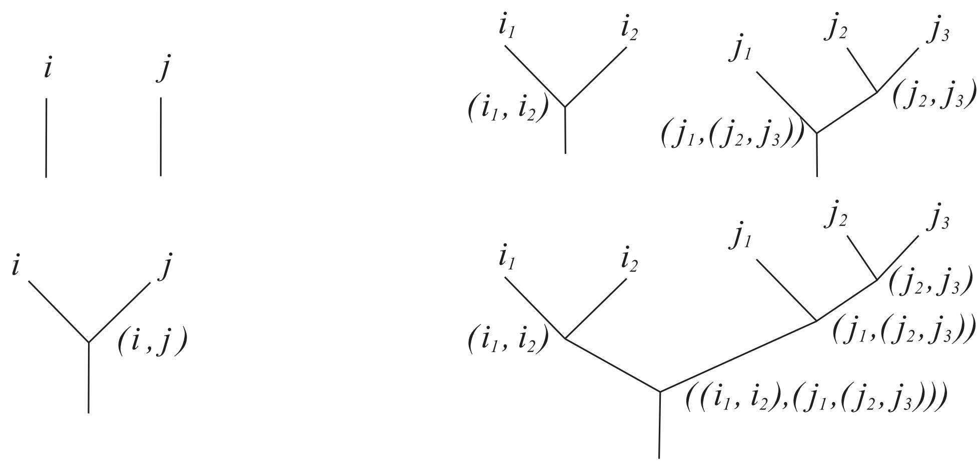

Figure 3: Rooted trees $t(i)$ and $t(j)$ (upper left) and their rooted product $t(i,j) = t(i)*t(j)$ (lower left); $t(i_1,i_2)$ and $t(j_1,(j_2,j_3))$ (upper right) and their rooted product $t((i_1,i_2),(j_1,(j_2,j_3))) = t(i_1,i_2)*t(j_1,(j_2,j_3))$ (lower right). In this figure all trivalent orientations are clockwise in the plane.

Thus, the non-root univalent vertices of the tree $t(I)$ associated to a bracket I are labelled by elements from the index set and the trivalent vertices correspond to sub-bracketings of I, with the trivalent vertex adjacent to the root corresponding to I.

Remark 8 The rooted product $*$ can be "realized" geometrically by a finger-move: Pushing a Whitney disk W_I through another Whitney disk W_J creates $W_{(I,J)}$ with $t(W_{(I,J)}) = t(W_I) * t(W_J)$.

This remark uses the upcoming assignment of a rooted tree $t(W)$ to a Whitney disk W inside a Whitney tower $\mathcal{W}$. In the easiest version, one starts with a root for W and then introduces one branching (trivalent vertex) while reading off which two sheets of $\mathcal{W}$ are paired by W. Then one continues with the same procedure for the two sheets to inductively obtain $t(W)$. In the next section we shall make this procedure precise, and in fact explain directly how orientations on the Whitney disks lead to vertex-orientations of the corresponding trees.

3.3 Rooted trees for oriented Whitney towers.

Let $\mathcal{W}$ be an oriented Whitney tower on order 0 surfaces A_i for $i = 1, 2, \ldots, m$. The orientations on the surfaces in $\mathcal{W}$ set up an indexing of the surfaces in $\mathcal{W}$ by bracketings I from $\{1, 2, \ldots, m\}$ and their corresponding rooted vertex oriented unitrivalent m-labelled trees $t(I)$ (3.2) via the following conventions:

A bracketing (i) of a singleton element i from the index set and the corresponding rooted order 0 tree $t(A_i) := t(i)$ are associated to each order 0 surface A_i. The bracket (I, J) and the corresponding tree $t(W_{(I,J)}) := t(I, J)$ are associated to a Whitney disk $W_{(I,J)}$, pairing intersections between W_I and W_J, with the ordering of the components I and J in the associated bracket (I, J) chosen so that the orientation of $W_{(I,J)}$ is the same as that given by orienting its boundary $\partial W_{(I,J)}$ from the negative intersection point to the positive intersection point *first* along W_I *then* back along W_J to the negative intersection point, together with a second inward pointing tangent vector.

We use brackets as subscripts to index surfaces in $\mathcal{W}$, writing A_i for an order 0 surface (dropping the brackets around the singleton i) and $W_{(i,j)}$ for a 1st order Whitney disk that pairs intersections between A_i and A_j, etc.. When writing $W_{(I,J)}$ for a Whitney disk pairing intersections between W_I and W_J, the understanding is that if a bracket I is just a singleton (i) then the surface $W_I = W_{(i)}$ is just the order 0 surface A_i. In general, the order of W_I is equal to the order of (i.e. the number of trivalent vertices of) $t(W_I)$.

It will be helpful to consider each tree $t(W_I)$ as a subset of $\mathcal{W}$: Assuming that $\mathcal{W}$ is based (Definition 7), map the vertices (other than the root) of $t(W_I)$ to the basepoints of the surfaces whose indices are contained as sub-brackets of I and map the edges (other than the edge adjacent to the root) of $t(W_I)$ to sheet-changing paths between basepoints, as illustrated in Figure 4 (disregarding, for the moment, the dotted loop which will be explained in 3.5). Then embed the root and its edge anywhere in the *negative corner* of W_I (see next paragraph).

It can be arranged that this mapping of $t(W_I)$ into $\mathcal{W}$ has the property that the trivalent orientations of $t(W_I)$ are induced by the orientations of the corresponding Whitney disks: Note that the pair of edges which pass from a trivalent vertex down into the lower order surfaces paired by a Whitney disk determine a "corner" of the Whitney disk which does not contain the other edge of the trivalent vertex. If this corner contains the *positive* intersection point paired by the Whitney disk, then the vertex orientation and the Whitney disk orientation agree. Our figures are drawn to satisfy this convention.

3.4 Orientation choices on Whitney disks.

Via our bracket-orientation convention, changing the orientation on a Whitney disk $W_{(I,J)}$ changes its tree from $t(W_{(I,J)}) = t(I,J)$ to $t(W_{(J,I)}) = t(J,I)$, i.e. changes the cyclic orientation of the associated trivalent vertex. In addition, changing the orientation of a *single* lower order Whitney disk W_K corresponding to a trivalent vertex of $t(W_{(I,J)})$ (so K is a sub-bracket of (I,J), with $K \neq (I,J)$) changes the cyclic orientations at exactly *two* trivalent vertices of $t(W_{(I,J)})$: the one corresponding to W_K and the adjacent one which corresponds to a Whitney disk pairing intersections between W_K and some other surface. That's because changing the orientation of W_K reverses the signs of the intersection points between W_K and anything else.

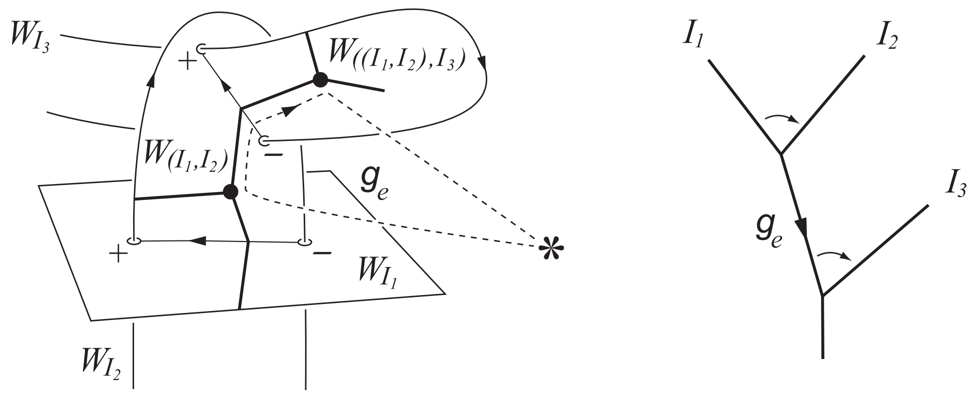

Figure 4: A Whitney disk $W_{((I_1,I_2),I_3)}$ and its associated tree $t(W_{((I_1,I_2),I_3)})$ shown (left) as a subset of the Whitney tower and (right) as an abstract rooted tree. The boundaries of the Whitney disks are oriented according to our bracket-orientation conventions using the indicated signs of the intersection points. The dashed path indicates a sheet-changing loop (based at the basepoint of the ambient 4–manifold X) which determines the element $g_e \in \pi_1 X$ decorating the corresponding oriented edge as described in 3.5.

3.5 Decorated trees for Whitney towers.

Let $t(W_I)$ be the (oriented labelled rooted) tree associated to a Whitney disk W_I in an oriented based Whitney tower $\mathcal{W}$ in a 4–manifold X. Thinking of $t(W_I)$ as a subset of $\mathcal{W}$ as described above, any edge e of $t(W_I)$, other than the root-edge, corresponds to a sheet-changing path connecting the basepoints of adjacent surfaces in $\mathcal{W}$. For a chosen orientation of e, this path together with the whiskers on the adjacent surfaces form an oriented loop which determines an element g_e of $\pi := \pi_1 X$ (Figure 4). Fixing (arbitrarily) orientations for all the (non-root) edges in $t(W_I)$ and labelling each oriented edge with an element of π in this way yields the *decorated rooted tree* associated to W_I (which will still be denoted by $t(W_I)$). Note that switching the orientation of e changes g_e to g_e^{-1} which explains the OR orientation reversal relation mentioned in 2.1 and shown in Figure 7. (Since we are working in an orientable 4–manifold, $\omega_1(g_e)$ is trivial.) Also, changing the choice of whisker on a Whitney disk has the effect of left multiplication on the group elements associated to the three edges adjacent to and oriented away from the trivalent vertex corresponding to the Whitney disk accounting for the HOL relation.

When decorations are understood, we will also denote a decorated tree by $t(I)$ where the underlying tree corresponds to the bracket I.

3.6 Decorated trees for intersection points.

If p is a transverse intersection point between W_I and W_J in $\mathcal{W}$ then the *decorated tree* t_p associated to p is defined as follows. Identify the roots of the decorated trees $t(W_I)$ and $t(W_J)$ to a single (non-vertex) point. The two edges that were adjacent to the roots of $t(W_I)$ and $t(W_J)$ now form a single edge e_p. Chose an orientation of e_p and decorate e_p by the element of π determined by the whiskers on W_I and W_J together with a path connecting the basepoints of W_I and W_J that changes sheets only at p with the orientation induced by e_p.

Thus, the decorated tree t_p is unrooted and every edge of t_p is oriented and decorated with an element of π. Note that the order of p is equal to the order of t_p (the number of trivalent vertices).

The mappings of $t(W_I)$ and $t(W_J)$ into $\mathcal{W}$ give rise to a mapping of t_p into $\mathcal{W}$: Just map the root vertices of W_I and W_J to p and the adjacent edges become a sheet-changing path between the basepoints of W_I and W_J (Figure 5). This mapping is an embedding of t_p into $\mathcal{W}$ if all the Whitney disks "beneath" W_I and W_J (corresponding to sub-brackets of I and J) are distinct.

We will sometimes keep track of the edge of t_p that corresponds to p by marking that edge with a small linking circle as in Figure 5; such a *punctured tree* will be denoted by t_p°.

It will be convenient to formalize the above description of the (unrooted) decorated tree t_p as a pairing (over the group π) of rooted decorated trees: Given

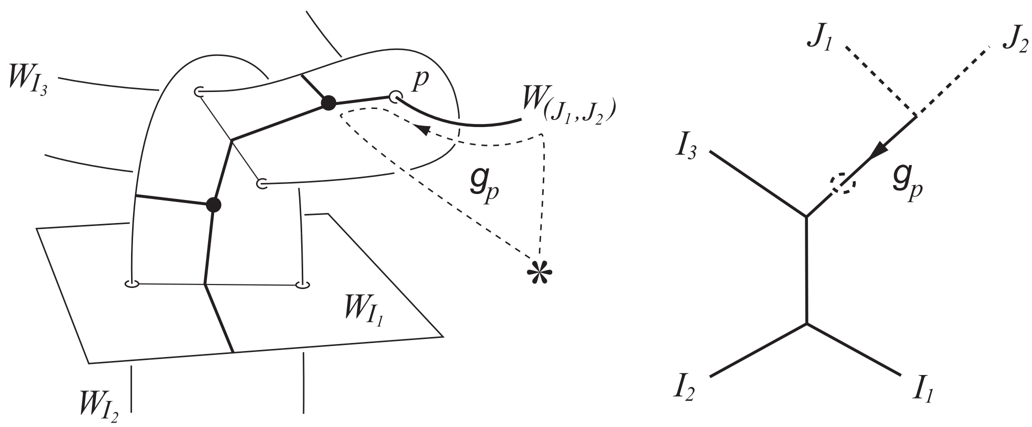

Figure 5: The punctured tree t_p° associated to an intersection point $p \in W_I \cap W_J$ (for $I = ((I_1, I_2), I_3)$ and $J = (J_1, J_2)$) shown as a subset of the Whitney tower and as an abstract labelled (punctured) tree. Decorations other than g_p are suppressed and the sheet-changing loop that determines g_p is indicated by the dashed path.

a pair $t(I)$ and $t(J)$ of rooted decorated trees and an element $g \in \pi$, define the *inner product* $t(I) \cdot_g t(J)$ to be the unrooted decorated tree gotten by identifying together the root vertices of $t(I)$ and $t(J)$ to a single (non-vertex) point in an edge labelled by g as illustrated in Figure 6. Thus, in this notation we have $t_p := t(W_I) \cdot_{g_p} t(W_J)$ for $p \in W_I \cap W_J$ as just described above.

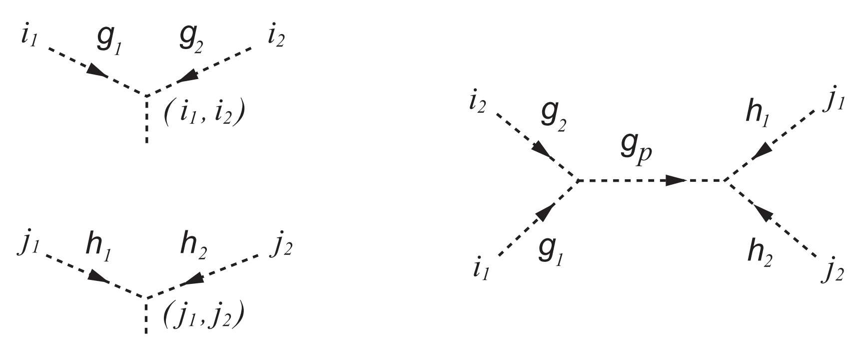

Figure 6: A pair of decorated rooted trees $t(I)$ and $t(J)$ corresponding to order 1 Whitney disks W_I and W_J with $I = (i_1, i_2)$ and $J = (j_1, j_2)$ (left), and the inner product $t_p = t(W_I) \cdot_{g_p} t(W_J) = t(I) \cdot_{g_p} t(J)$ associated to an order 2 intersection point $p \in W_I \cap W_J$ (right).

3.7 The antisymmetry AS relation.

If a Whitney tower $\mathcal{W}$ is oriented then there is one more piece of information that we need to keep track of: the sign ϵ_p of an unpaired intersection point

$$p \in W_I \cap W_J \subset \mathcal{W}.$$

ϵ_p is computed, in the usual way, by comparing the orientation determined by W_I and W_J at p with the orientation of the ambient 4–manifold X at p.

Changing the orientation on the Whitney disk W_I changes the *signed tree* $\epsilon_p \cdot t_p$ by the AS antisymmetry relation mentioned in 2.1: The cyclic orientation of the vertex corresponding to W_I in t_p is switched and so is the sign ϵ_p of the intersection with W_J. Moreover, changing the orientation of a single Whitney disk, other than W_I or W_J, preserves the sign ϵ_p and changes the cyclic orientations at *two* trivalent vertices of t_p, as pointed out above in Section 3.4. Consequently, working modulo the AS relation makes the signed tree $\epsilon_p \cdot t_p$ independent of the choices of orientations for the Whitney disks in $\mathcal{W}$.

The dependence on orientations for the original sheets A_i remains: changing the orientation of one A_i introduces an additional sign into $\epsilon_p \cdot t_p$ if t_p has an odd number of i-labelled vertices.

3.8 The intersection tree $\tau_n(\mathcal{W})$.

We would next like to add up the unpaired intersection points of a given Whitney tower in some algebraic structure. For that purpose, let $\mathcal{T}_n(\pi, m)$ denote the abelian group generated by (isomorphism classes of) decorated trees of order n modulo the relations shown in Figure 7. That is, each generator is an (unrooted) unitrivalent tree having

- n cyclically oriented trivalent vertices,
- $n+2$ univalent vertices labelled by elements of $\{1, \dots, m\}$, and
- $2n+1$ oriented edges decorated by elements of π.

Definition 9 Let $\mathcal{W}$ be an order n Whitney tower on properly immersed simply-connected oriented surfaces $A_1, \dots, A_m$ in a 4–manifold X. (In fact, the A_i only need to be π_1-null, see [11].) Define the n-th order *intersection tree* of $\mathcal{W}$ by

$$\tau_n(\mathcal{W}) := \sum_p \epsilon_p \cdot t_p \quad \in \quad \mathcal{T}_n(\pi, m)$$

where the sum is over all order n intersection points p in $\mathcal{W}$.

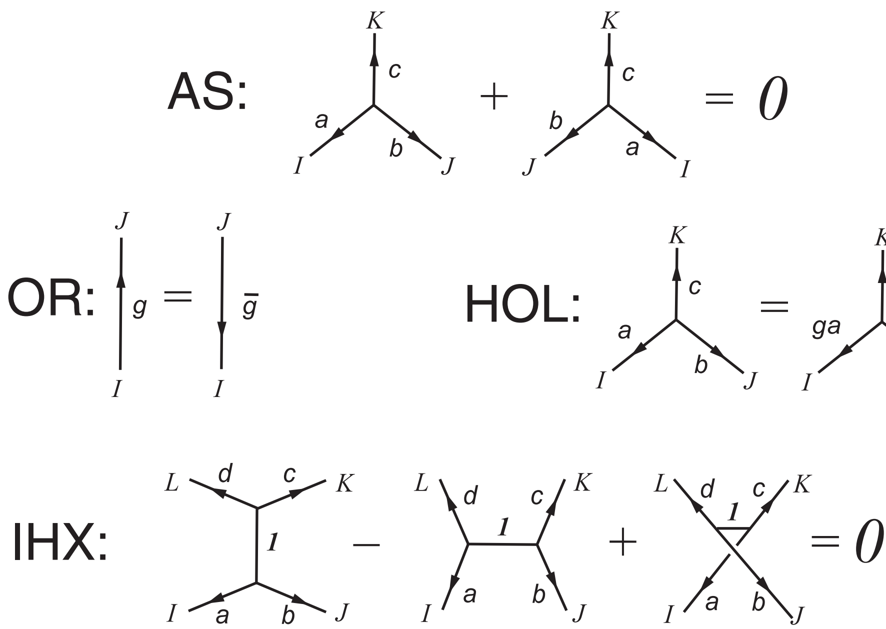

Figure 7: The AS, OR, HOL and IHX relations in $\mathcal{T}_n(\pi, m)$ for a, b, c, d, 1 and g in π with $\overline{g} = g^{-1}$. All trivalent orientations are induced from a fixed orientation of the plane.

As explained above, the AS relations make sure that $\tau_n(\mathcal{W})$ actually does not depend on the choice of orientations for the Whitney disks. Similarly, the HOL and OR relations make sure that $\tau_n(\mathcal{W})$ does not depend on the choice of whiskers, or edge orientations. In other words, $\tau_n(\mathcal{W})$ is defined by first choosing whiskers and orientations (on edges and Whitney disks) and then proving independence of these choices.

Remark 10 Using the HOL relation or, more concretely, by choosing the whiskers on the Whitney disks appropriately, one can normalize the trees t_p so that all interior edges and one univalent edge are decorated with the trivial group element $1 \in \pi$. Thus, one can interpret $\tau_n(\mathcal{W})$ as living in a quotient of the integral group ring of the $(n+1)$-fold product of π.

By slightly refining our notation, signs can be associated formally to all tree edges and the edge decorations can be extended linearly to elements of the group ring $\mathbb{Z}[\pi]$ (compare [14], [15]). Similarly, one can extend the labels on the univalent vertices to the free abelian group on $\{1, \ldots, m\}$.

4 Proof of Theorem 2

Our proof of Theorem 2 will be constructive in the sense that we describe how to build the next order Whitney tower by geometrically realizing all the relations in $\mathcal{T}_n(\pi, m)$. However, it should be mentioned that since the groups $\mathcal{T}_n(\pi, m)$ do not in general have a canonical basis we are sidestepping the "word problem" in $\mathcal{T}_n(\pi, m)$. The main construction (Lemma 15) of the proof shows how to exchange *algebraic cancellation* of pairs of intersection points for *geometric cancellation* (by Whitney disks) in the case that the intersection points are *simple* (have certain standard right- or left-normed trees, 4.5). This algebraic cancellation occurs in the lift $\widehat{\mathcal{T}}$ of $\mathcal{T}$ which forgets the IHX relation. The general case is then reduced to this case using geometric IHX constructions from [9] and [28] to show that an order n Whitney tower $\mathcal{W}$ with $\tau_n(\mathcal{W}) = 0 \in \mathcal{T}_n(\pi, m)$ can be modified so that all order n intersections come in simple algebraically cancelling pairs.

To simplify the exposition and highlight the combinatorial structure of Whitney towers, we will emphasize the simply-connected case, often dropping the group π from notation. Refining the constructions to cover the general case for the most part only requires checking that whiskers can be (re)-chosen appropriately. At a first reading it doesn't hurt to ignore group elements entirely and only the simply-connected version of Theorem 2 will be used later in the proof of Theorem 4.

We begin with some notation and lemmas. All Whitney towers are assumed oriented, labelled and based.

4.1 Geometric intersection trees for Whitney towers

For an (oriented, labelled, based) Whitney tower $\mathcal{W}$ define $t_n(\mathcal{W})$, the (n-th order, oriented) *geometric intersection tree* of $\mathcal{W}$, to be the disjoint union of signed (decorated) trees

$$t_n(\mathcal{W}) := \amalg_p \epsilon_p \cdot t_p$$

over all unpaired order n intersection points $p \in \mathcal{W}$. (An unsigned version of $t_n(\mathcal{W})$ was defined for unoriented Whitney towers in [28].) The next two pairs of definitions and lemmas will illustrate how $t_n(\mathcal{W})$ captures the essential geometric structure of $\mathcal{W}$.

4.2 Split subtowers

The Whitney disks in an arbitrary Whitney tower may have multiple self-intersections and intersections with other surfaces. However, it is not difficult to modify an arbitrary Whitney tower so that each Whitney disk is embedded and contains either a single Whitney arc or unpaired intersection point (Lemma 13 below). This is best expressed using the notion of *split subtowers* and splitting a Whitney tower into split subtowers will serve to simplify geometric constructions and combinatorial arguments.

The purpose of constructing a Whitney tower is to provide information on the homotopy classes (rel boundary) of its order 0 surfaces. However, when describing and manipulating *subsets* of a Whitney tower it is natural to consider *sub*towers on sheets of surfaces which are not *properly* immersed:

Definition 11 A *subtower* is a Whitney tower except that the boundaries of the immersed order 0 surfaces in a subtower are allowed to lie in the interior of the 4–manifold (instead of being required to lie in the boundary). The boundaries of the order 0 surfaces in a subtower are still required to be embedded. The notions of *order* for intersection points and Whitney disks are the same as in Definition 7.

In this paper we will only be concerned with subtowers whose order 0 surfaces are sheets in the order 0 surfaces of an actual Whitney tower. In this case, the surfaces of the subtower inherit the same orientations and indexing by brackets as the Whitney tower. Thus, the association of decorated trees to surfaces and intersection points is also the same.

Definition 12 A subtower $\mathcal{W}_p$ is *split* if it satisfies all of the following:

(i) $\mathcal{W}_p$ contains a single unpaired intersection point p,

(ii) the order 0 surfaces of $\mathcal{W}_p$ are all embedded 2–disks,

(iii) the Whitney disks of $\mathcal{W}_p$ are all embedded,

(iv) the interior of any surface in $\mathcal{W}_p$ either contains p or contains a single Whitney arc of a Whitney disk in $\mathcal{W}_p$,

(v) $\mathcal{W}_p$ is connected (as a 2–complex in the 4–manifold).

Moreover, a Whitney tower $\mathcal{W}$ is called *split* if all the unpaired intersection points of $\mathcal{W}$ are contained in disjoint split subtowers on sheets of the order 0 surfaces of $\mathcal{W}$.

Note that a normal thickening of a split subtower $\mathcal{W}_p$ in the ambient 4–manifold is just the 4–disk D^4 which is a regular neighborhood of the embedded tree t_p associated to the unpaired intersection point p.

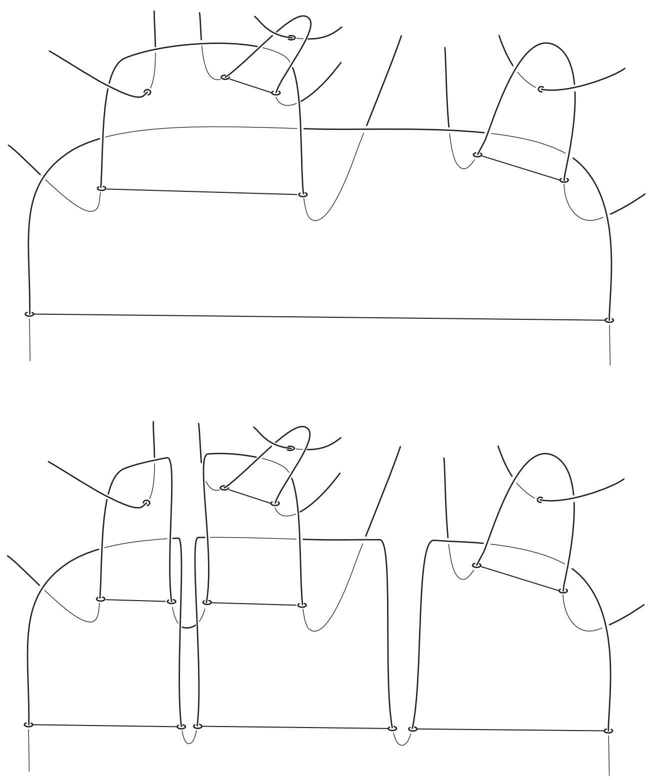

Figure 8: Part of a Whitney tower before (top) and after (bottom) applying the splitting procedure described in the proof of Lemma 13.

4.3 Split Whitney towers

The splitting of a Whitney tower into split subtowers described in the following lemma is analogous to Krushkal's splitting of a grope into genus one gropes [22].

Lemma 13 *Let $\mathcal{W}$ be a Whitney tower on order 0 surfaces A_i. Then there exists a split Whitney tower $\mathcal{W}_{\text{split}}$ contained in any regular neighborhood of $\mathcal{W}$ such that:*

(i) The order 0 surfaces A_i' of $\mathcal{W}_{\text{split}}$ only differ from the A_i by finger moves.

(ii) The geometric intersection trees $t(\mathcal{W})$ and $t(\mathcal{W}_{\text{split}})$ are isomorphic.

The isomorphism in item (ii) includes decorations and signs.

Proof Starting with the highest order Whitney disks of $\mathcal{W}$, apply finger moves as indicated in Figure 8. Working down through the lower order Whitney disks yields the desired $\mathcal{W}_{\text{split}}$. Choosing whiskers and orientations appropriately for the new Whitney disks preserves the decorations on the trees associated to the unpaired intersection points. □

An advantage of splitting a Whitney tower is that the geometric intersection tree sits as an *embedded* subset (3.6) and all the singularities of the split Whitney tower are contained in disjointly embedded 4–balls, each of which is a regular neighborhood of an intersection point tree. In this sense the decomposition of a Whitney tower into split subtowers corresponds to the idea that the trees associated to the unpaired intersection points capture the essential structure of a Whitney tower. The next lemma can be interpreted as justifying that this essential structure is indeed captured by the *un*-punctured trees rather than the punctured trees in the sense that an unpaired intersection point (corresponding to a punctured edge) can be "moved" to any other edge of its tree.

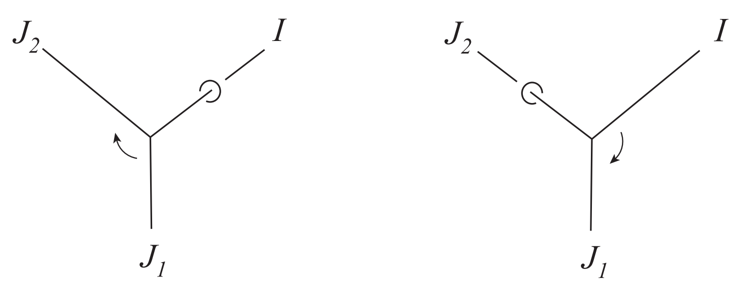

Figure 9: A local picture of the tree associated to the split subtower $\mathcal{W}$ before (left) and $\mathcal{W}'$ after (right) the Whitney move in the proof of Lemma 14 illustrated in Figure 10 and Figure 11.

Lemma 14 *Let $\mathcal{W} \subset X$ be a split subtower on order 0 sheets s_i with unpaired intersection point $p = W_I \cap W_J \subset \mathcal{W}$. Denote by $\nu(\mathcal{W})$ a normal thickening of $\mathcal{W}$ in X so that $\partial s_i \subset \partial\nu(\mathcal{W}) \subset \nu(\mathcal{W}) \cong D^4$. If I' and J' are any brackets such that the decorated trees $t(I') \cdot t(J') = t_p = t(I) \cdot t(J)$, then after a homotopy* *(rel ∂) of the s_i in $\nu(\mathcal{W})$ the s_i admit a split subtower $\mathcal{W}' \subset \nu(\mathcal{W})$ with single unpaired intersection point $p' = W_{I'} \cap W_{J'} \subset \mathcal{W}'$ such that $\epsilon_{p'} \cdot t_{p'} = \epsilon_p \cdot t_p$.*

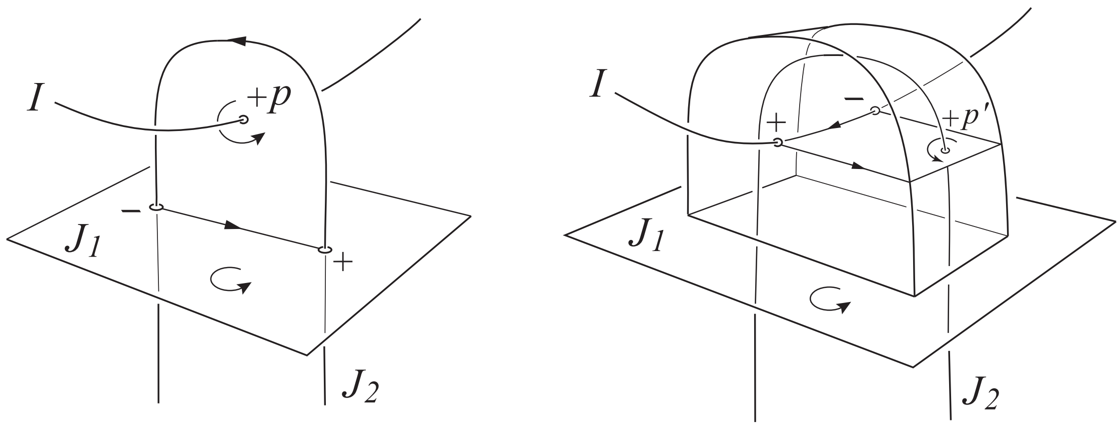

Figure 10: The unpaired intersection point $p = W_I \cap W_J$ in the split subtower $\mathcal{W}$ of Lemma 14 (left), and the unpaired intersection point $p' = W_{I'} \cap W_{J'}$ in $\mathcal{W}'$ after the Whitney move (right). Signs and orientations are indicated for the case $\epsilon_p = +$, with brackets corresponding to the trivalent orientations in Figure 9.

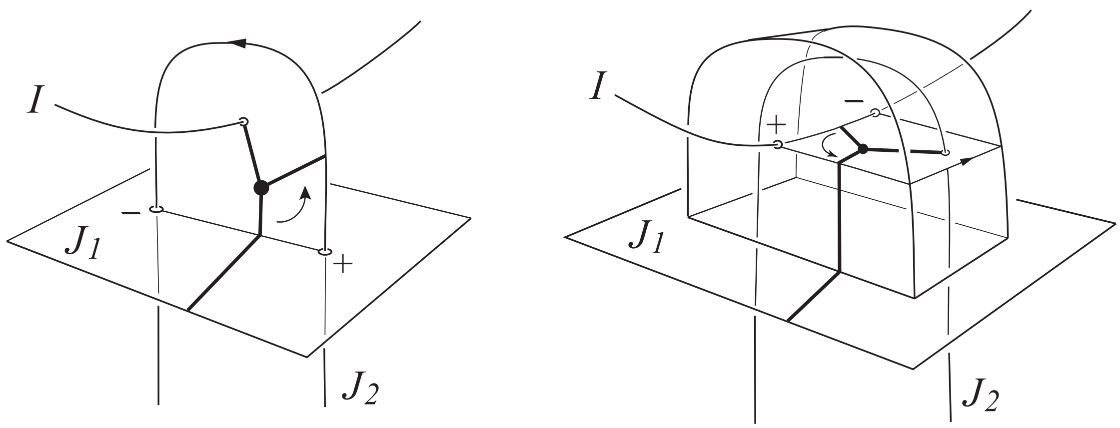

Figure 11: This figure shows that the oriented punctured trees associated to p and p' in Figure 10 differ as indicated in Figure 9.

Proof (of Lemma 14) It is enough to show that the puncture in t_p° can be "moved" to either *adjacent* edge, since by iterating it can be moved to any edge of t_p. Specifically, it is enough to consider the case where $J = (J_1, J_2)$, $I' = (I, J_1)$ and $J' = J_2$ so that $I \cdot (J_1, J_2) = (I, J_1) \cdot J_2$ as in Figure 9. (Here we are assuming that W_J is not order 0 since if both W_I and W_J are order 0 there is nothing to prove.) The proof is given by the maneuver illustrated in Figure 10: Use the Whitney disk W_J to guide a Whitney move on W_{J_1}. This eliminates the intersections between W_{J_1} and W_{J_2} (as well as eliminating W_J and p) at the cost of creating a new cancelling pair of intersections between W_{J_1} and W_I. This new cancelling pair can be paired by a Whitney disk $W_{(I,J_1)}$ having a single intersection point p' with W_{J_2}. That this achieves the desired effect on the punctured tree can be seen in Figure 11 by referring to the signs and orientations in Figure 10. See also the discussion in pages 20–22 of [30] which includes group elements. □

4.4 Algebraically and geometrically cancelling pairs

Let $\widehat{\mathcal{T}}_n(\pi,m)$ denote the group of order n decorated trees modulo all the relations in Figure 7 *except* the IHX relation. We say that a pair of intersection points p_+ and p_- in $\mathcal{W}$ *cancel algebraically* if $\epsilon_{p_+}\cdot t_{p_+} = -\epsilon_{p_-}\cdot t_{p_-} \in \widehat{\mathcal{T}}_n(\pi,m)$. There is a summation map that sends the disjoint union $t_n(\mathcal{W}) = \amalg_p \epsilon_p \cdot t_p$ to an element $\widehat{\tau}_n(\mathcal{W}) := \sum_p \epsilon_p \cdot t_p \in \widehat{\mathcal{T}}_n(\pi,m)$ and the vanishing of $\widehat{\tau}_n(\mathcal{W})$ is equivalent to being able to arrange all of the order n intersection points of $\mathcal{W}$ into algebraically cancelling pairs.

Given an algebraically cancelling pair $p_\pm$ in a split Whitney tower, one can chose orientations and whiskers on the Whitney disks in the split subtowers containing $p_\pm$ so that the trees $t_{p_\pm}$ have identical orientations (and decorations) with $\epsilon_{p_+} = -\epsilon_{p_-}$. (This is because the OR, HOL and AS relations are realized by these choices, as described in Sections 3.5 and 3.7.)

A pair of intersection points p_+ and p_- in $\mathcal{W}$ *cancel geometrically* if they can be paired by a Whitney disk. Geometric cancellation implies algebraic cancellation, but the converse is not true since two algebraically cancelling intersection points might not lie on the same Whitney disks.

The next lemma gives sufficient conditions for a sort of converse involving some additional work.

4.5 Simple intersection points and the transfer lemma

Following the terminology of [19] for iterated commutators of group elements, we say that an intersection point $p \in \mathcal{W}$ is *simple* if its tree t_p is simple (right- or left-normed) as illustrated in Figure 12. The proof of the next lemma shows

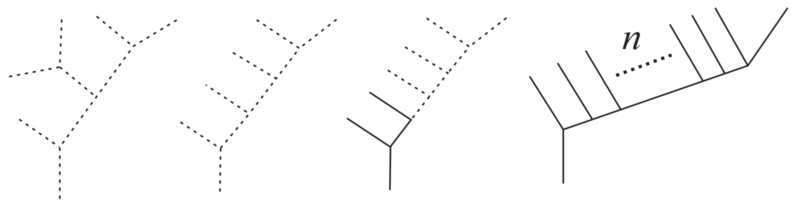

Figure 12: From left to right, the non-simple tree of lowest order (order 4) and the simple trees of order 4, 5 and $6+n$.

how to exchange simple algebraically cancelling pairs of intersection points for geometrically cancelling pairs.

Lemma 15 *Let $\mathcal{W}$ be an order n Whitney tower on order 0 surfaces A_i such that all order n intersection points of $\mathcal{W}$ come in* simple *algebraically cancelling pairs. Then the A_i are homotopic (rel boundary) to A_i' which admit an order $(n+1)$ Whitney tower.*

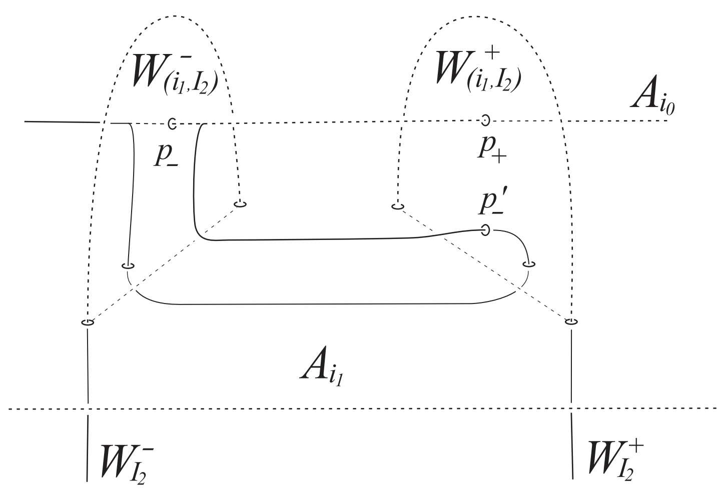

Figure 13

Proof We will describe a modification of $\mathcal{W}$ which exchanges one algebraically cancelling simple pair of order n for another at the cost of only creating geometrically cancelling pairs. Iterating this modification will, at the n-th iteration, exchange an algebraically cancelling pair for *only* geometrically cancelling pairs. This modification is described in [36] for the case $n=1$ in a simply-connected manifold. (See also [30] for the $n=1$ non-simply-connected case.) Applying this procedure to all algebraically cancelling pairs will complete the proof. We will discuss only the simply-connected case; the reader can easily add group elements to the figures (as in [30]).

We may assume that $\mathcal{W}$ is split by Lemma 13. Let p_+ and p_- be a simple algebraically cancelling pair of order n intersection points in $\mathcal{W}$. By "pushing the puncture out to an end of the simple tree" using Lemma 14, we may further assume that p_+ and p_- are intersections between some order 0 surface A_{i_0} and order n Whitney disks $W^+_{I_1}$ and $W^-_{I_1}$ respectively where, for this proof only, I_k will denote a simple bracket of the form $I_k := (i_k, (i_{(k+1)}, (\ldots, (i_n, i_{(n+1)})\ldots)) = (i_k, I_{(k+1)})$ for $1 \leq k \leq n+1$ and $I_{(n+1)} = i_{(n+1)}$. The first step in the modifi-

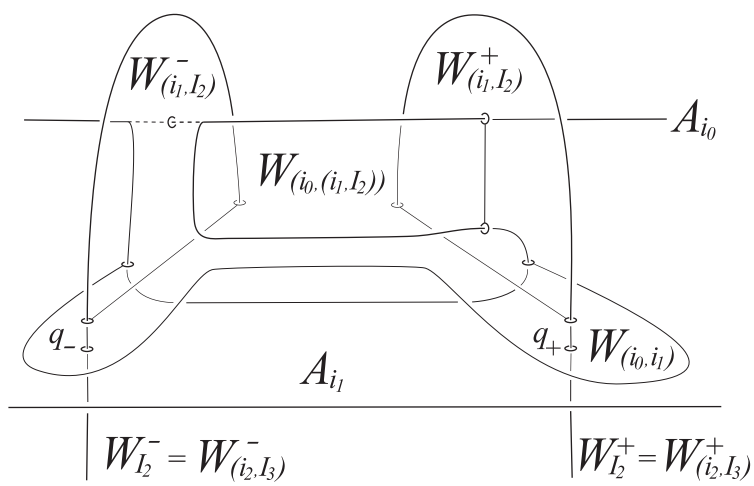

Figure 14

cation is illustrated in Figure 13 which shows how to exchange $p_- \in A_{i_0} \cap W^-_{I_1}$ for $p'_- \in A_{i_0} \cap W^+_{I_1}$, which cancels geometrically with p_+, at the cost of creating a geometrically cancelling pair of intersection points between A_{i_0} and A_{i_1}. Note that this first step is possible because both A_{i_0} and A_{i_1} are *connected*. The modification is completed by choosing Whitney disks for the new geometrically cancelling pairs as illustrated in Figure 14, which shows that a new algebraically cancelling pair $q_\pm \in W_{(i_0,i_1)} \cap W^\pm_{I_2}$ has been created (recall that boundaries of Whitney disks must be disjointly embedded). In the case $n=1$, $q_\pm$ would also cancel geometrically since then $I_{(n+1)} = i_{(n+1)}$ means that $W^+_{I_2} = W^-_{I_2} = A_{i_2}$ which is connected. Note that $W_{(i_0,i_1)}$ is embedded (in a neighborhood of a contractible 1–complex) and contains only the pair $q_\pm$ in its interior. The Whitney disk $W_{(i_0,(i_1,I_2))}$ may intersect anything but we don't care because it is a Whitney disk of order $n+1$ and hence can only contain intersections of order strictly greater than n. Now, assuming $n \geq 2$, apply this

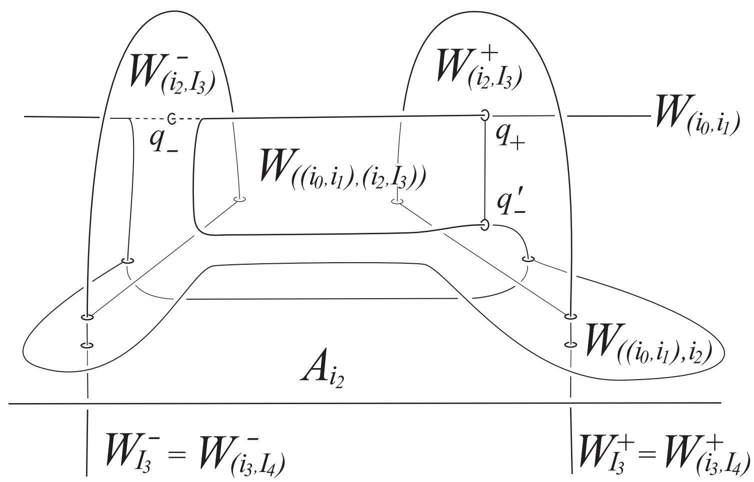

Figure 15

modification to $q_\pm$ as illustrated in Figure 15. Note that this is only possible because we have the *connected* surface A_{i_2} to "push along", since we originally started with the *simple* pair $p_\pm$ so that $W^{\pm}_{I_2} = W_{(i_2,I_3)}$. The k-th iteration

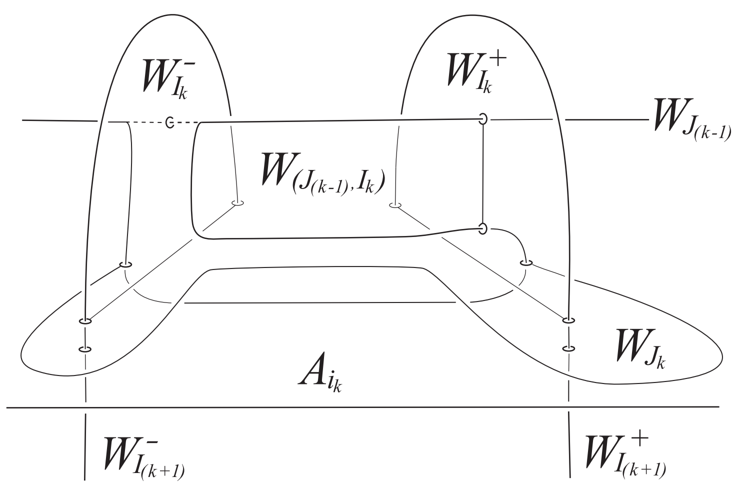

Figure 16

of this modification is illustrated in Figure 16 where, for this proof only, we denote the simple bracket $J_k := (\ldots((i_0, i_1), i_2), \ldots, i_k)$ for $1 \leq k \leq n$. The procedure terminates when $k = n$ meaning that $W^{\pm}_{I_{(k+1)}} = W_{I_{(n+1)}} = A_{i_{(n+1)}}$ which is connected so only geometrically cancelling pairs are created.

This procedure can be applied to all the (simple) algebraically cancelling pairs: One can always find disjoint arcs between Whitney arcs in the A_{i_k} to guide the modification and all new Whitney disks of order $\leq n$ are contained in neighborhoods of these arcs so that no unexpected intersections of order less than or equal to n are created. □

4.6 Geometric IHX and the Proof of Theorem 2

Given $\mathcal{W}$ as in Theorem 2, we will reduce the proof to the case handled by Lemma 15 by using geometric constructions and results from [9] and [28]. Achieving the hypotheses of Lemma 15 will involve two steps: First $\mathcal{W}$ will be modified to have only algebraically cancelling pairs by using the "4–dimensional IHX construction" in [9]. Then the algebraically cancelling pairs will be exchanged for simple algebraically cancelling pairs, using a related IHX construction of [28]. This second step is based on the effect of doing a Whitney move on a Whitney disk in a split subtower and mimics the usual algebraic proof that the group of unitrivalent trees modulo the IHX and AS relations is spanned by simple trees ([1], [7]).

4.7 Creating algebraically cancelling pairs.

The vanishing of $\tau_n(\mathcal{W}) \in \mathcal{T}^t_n(\pi, m)$ means that $\tau_n(\mathcal{W})$ lifts to $\widehat{\tau}_n(\mathcal{W}) \in \operatorname{span}\{\mathrm{I} - \mathrm{H} + \mathrm{X}\} < \widehat{\mathcal{T}}_n(\pi, m)$. To get only algebraically cancelling pairs we apply the following corollary of the 4*–dimensional IHX Theorem* in [9]:

Proposition 16 *Let $\mathcal{W}$ be any order n Whitney tower on order 0 surfaces A_i. Then, given any decorated order n unitrivalent trees* I, H *and* X *differing only by the local* IHX *relation of Figure 7, there exists an order n Whitney tower $\mathcal{W}'$ on A'_i homotopic (rel boundary) to the A_i such that*

$$t_n(\mathcal{W}') = t_n(\mathcal{W}) + \mathrm{I} - \mathrm{H} + \mathrm{X}.$$

Note that the "sum" on the right hand side is really a disjoint union of signed decorated trees; the summation map takes this equation to the corresponding equation in $\widehat{\mathcal{T}}_n(\pi, m)$.

Proof As observed in Remark 8, creating a "clean" Whitney disk by applying a finger move to surfaces in a Whitney tower "realizes" the rooted product $*$ on the corresponding rooted trees. Since finger moves are supported near arcs, one can modify $\mathcal{W}$ to create any number of clean Whitney disks realizing arbitrary rooted decorated trees without changing $t_n(\mathcal{W})$. Let W^i, $i = 1, 2, 3, 4$ be four such Whitney disks which correspond to the four fixed vertices of the trees I, H and X in the statement. (Of course if any of the fixed vertices is univalent then the corresponding "Whitney disk" is just an order 0 surface.)

Now the 4–dimensional IHX Theorem of [9] says that there exists an order 2 Whitney tower $\mathcal{W}_{\text{IHX}}$ on oriented 2–spheres A_i, $i = 1, 2, 3, 4$, in a 4–ball having geometric intersection tree $t_2(\mathcal{W}_{\text{IHX}})$ equal precisely to the order 2 IHX relation. So by tubing A_i into W^i, for each i, we can get $\mathcal{W}'$ as desired. No unexpected intersections are created since the entire construction takes place near a collection of arcs and the (arbitrarily small) 4–ball. (In the decorated case the desired group elements are controlled by the tubes.) $\square$

So by applying Proposition 16 as necessary we can assume that $\widehat{\tau}_n(\mathcal{W}) = 0 \in \widehat{\mathcal{T}}_n(\pi, m)$ which means that all order n intersection points can be arranged in algebraically cancelling pairs.

4.8 Simplifying the cancelling pairs.

In case there are algebraically cancelling pairs which are not simple, we appeal to results in [28]: Proposition 7.1 of [28] describes an algorithm for modifying a Whitney tower to have only simple intersection points. This geometric algorithm, which mimics the algebraic algorithm described in [1] and [7], depends on a "Whitney move" version of the IHX relation (Lemma 7.2 of [28]) which replaces a split subtower $\mathcal{W}_p$ by two split subtowers $\mathcal{W}_{p'}$ and $\mathcal{W}_{p''}$ and has the effect of replacing $\epsilon_p \cdot t_p = \text{I}$ by $\text{H} - \text{X} = \epsilon_{p'} \cdot t_{p'} + \epsilon_{p''} \cdot t_{p''}$ in the geometric intersection tree. The point of the algorithm is that the trees H and X are "closer" to being simple and by iterating one is eventually left with only simple trees. (The construction is supported in a neighborhood of $\mathcal{W}_p$ so no unwanted intersections are created.) Although Proposition 7.1 and Lemma 7.2 of [28] are only proved in the unoriented undecorated case it is not hard to add signs to the intersection points in the diagrams in [28] and apply the conventions of this paper, especially having seen the related proof of Lemma 15 above.

So in the present setting we have only algebraically cancelling pairs of order n intersection points in an order n Whitney tower $\mathcal{W}$ which we may assume is split by Lemma 13. If any of these cancelling pairs are not simple, then we apply the just mentioned IHX algorithm of [28] *pairwise* (so as to preserve $\widehat{\tau}_n(\mathcal{W}) = 0 \in \widehat{\mathcal{T}}_n(\pi, m)$) until we are left with only simple algebraically cancelling pairs. The proof of Theorem 2 is now complete by Lemma 15. $\square$

5 Proof of Theorem 4

The proof of Theorem 4 uses results from [17], [23] and [28] as well Theorem 2 to compare an arbitrary link L to certain well-known standard links which generate the first non-vanishing Milnor and Z^t invariants.

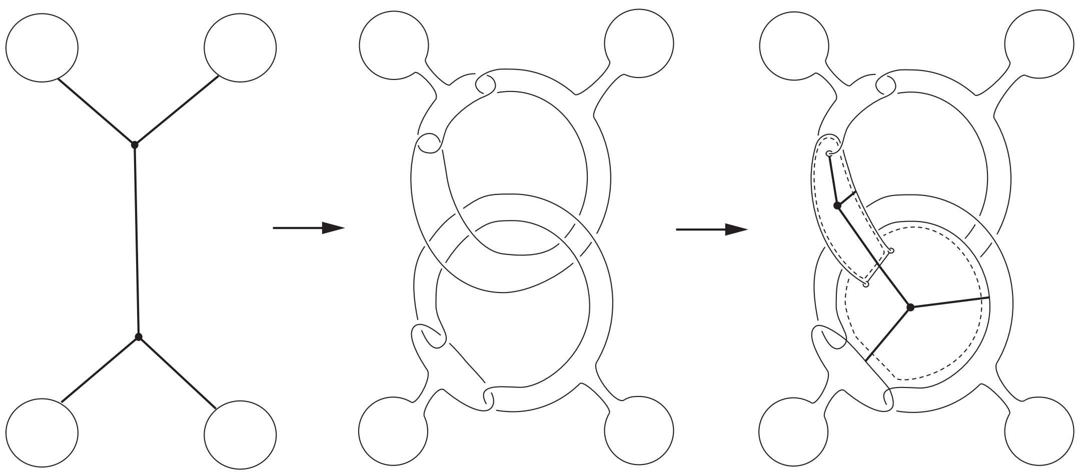

Figure 17: From left to right: An order 2 (positively signed) vertex-oriented tree I whose univalent vertices correspond to the components of an unlink; The Bing-Cochran-Habiro link L_{I}; An order 2 Whitney tower $\mathcal{W}$ bounded by L_{I} with $\tau_2(\mathcal{W}) = \mathrm{I}$.

5.1 Bing-Cochran-Habiro links

Given a collection σ of signed labelled vertex-oriented order n trees, Cochran [4] and Habiro [16] have described, using Bing doubling and Clasper surgery respectively, how to construct (from the unlink) a link L_σ such that $K_n(L_\sigma)$ is represented by σ (considered as a sum). Habiro's construction applies more generally to unitrivalent *graphs*, but for trees the two constructions coincide (by applying Kirby calculus to a framed link surgery description).

Given such a *Bing-Cochran-Habiro link* L_σ, we will use the following two facts:

(i) L_σ bounds an order n Whitney tower $\mathcal{W}_\sigma$ with $\tau_n(\mathcal{W}_\sigma) = \sigma \in \mathcal{T}_n(m)$.

(ii) $K_n(L_\sigma) = \sigma \in \mathcal{T}_n(m) \otimes \mathbb{Q}$.

The Whitney tower $\mathcal{W}$ in statement (i) is easily constructed by "pulling apart" a Bing double in Cochran's construction (see Figure 17): This creates Whitney disks whose boundaries are essentially the *derived links* in [4] and each $t_p \in \sigma$ corresponds to a *derived linking*. Alternatively, starting with Habiro's clasper surgery description one can apply the translation to *grope cobordism* of [7] and then the translation to Whitney towers of [28] and [9].

For statement (ii), see Section 8 of [17]. Although [17] works with *string* links, the first non-vanishing term of $Z^t(L)-1$ is equal to the first non-vanishing term of $Z^t(SL) - 1$ where SL is any string link whose closure is L (see Section 5 of [27]).

5.2 Whitney towers and the Kontsevich integral

Let L and $\mathcal{W}$ be as hypothesized in Theorem 4. Denote by σ any disjoint union of signed (labelled vertex-oriented) trees which represents $\tau_n(\mathcal{W}) \in \mathcal{T}_n(m)$, e.g. the geometric intersection tree $t(\mathcal{W})$ of $\mathcal{W}$ (4.1). Let L_σ be a Bing-Cochran–Habiro link formed from the unlink using σ. Then, by (i) of 5.1, L_σ bounds an order n Whitney tower $\mathcal{W}_\sigma$ in D^4 with $\tau_n(\mathcal{W}_\sigma) = \tau_n(\mathcal{W}) \in \mathcal{T}_n(m)$. Now think of $\mathcal{W}$ and $\mathcal{W}_\sigma$ as each sitting in a copy of $S^3 \times I$ (D^4 with a neighborhood of a point removed). By gluing together the two copies of $S^3 \times I$ (along the S^3 boundary of the removed neighborhoods) and connecting each order zero 2–disk of $\mathcal{W}$ with the corresponding order zero 2–disk of $\mathcal{W}_\sigma$ by a small tube we get properly immersed annuli A_i in $S^3 \times I$ cobounded by the link components. Since the tubes may be chosen to avoid creating new intersection points, the A_i admit an order n Whitney tower $\mathcal{W}'$ with

$$\tau_n(\mathcal{W}') = \tau_n(\mathcal{W}) - \tau_n(\mathcal{W}_\sigma) = 0 \in \mathcal{T}_n(m)$$

where the minus sign comes from reversing the orientation of one of the two copies of $S^3 \times I$. By Theorem 2, the vanishing of $\tau_n(\mathcal{W}')$ implies that (after a homotopy rel boundary) the A_i admit a Whitney tower of order n, that is, L and L_σ are *order n Whitney equivalent*. By the main theorem in [28], order n Whitney equivalence implies (in fact is equivalent to) *class* $(n+1)$ *grope concordance*, meaning that we can conclude that the components of L and L_σ cobound disjoint properly embedded annulus-like gropes of class $(n+1)$. This implies, by [23] Corollary 4.2, that L and L_σ have the same μ-invariants of length less than or equal to $(n+1)$. It follows from [17] that $K_n(L) = K_n(L_\sigma)$ which is equal to $\sigma \in \mathcal{T}_n(m) \otimes \mathbb{Q}$ by (ii) of 5.1 above. □

References

[1] **D Bar-Natan**, *Vassiliev homotopy string link invariants*, J. Knot Theory Ramifications 4 (1995) 13–32.

[2] **A Casson**, *Three Lectures on new infinite constructions in 4–dimensional manifolds*, Notes prepared by L Guillou, Prepublications Orsay 81T06 (1974).

[3] **A Casson**, *Link cobordism and Milnor's invariant*, Bull. London Math. Soc. 7 (1975) 39–40.

[4] **T Cochran**, *Derivatives of links, Milnor's concordance invariants and Massey products*, Mem. Amer. Math. Soc. Vol. 84 No. 427 (1990).

[5] **T Cochran**, **K Orr**, **P Teichner**, *Knot concordance, Whitney towers and L^2-signatures*, Annals of Math., Volume 157 (2003) 433–519.

[6] **T Cochran**, **P Teichner**, *Knot concordance and von Neumann ρ-invariants*, Preprint 2004.

[7] **J Conant**, **P Teichner**, *Grope cobordism of classical knots*, Topology 43 (2004) 119–156.

[8] **J Conant**, **P Teichner**, *Grope Cobordism and Feynman Diagrams*, to appear in Math. Annalen.

[9] **J Conant**, **R Schneiderman**, **P Teichner**, *A geometric IHX relation in 3 and 4 dimensions*, In preparation.

[10] **M Freedman**, *The topology of four-dimensional manifolds*, J. Diff. Geom. 17 (1982) 357–453.

[11] **M Freedman**, **F Quinn**, *The topology of 4–manifolds*, Princeton Math. Series 39 Princeton, NJ, (1990).

[12] **C H. Giffen** *Link concordance implies link homotopy*, Math. Scand. 45 (1979) 243–254.

[13] **D Goldsmith** *Concordance implies homotopy for classical links in M^3*, Comment. Math. Helvetici 54 (1979) 347–355.

[14] **S Garoufalidis**, **A Kricker**, *A rational noncommutative invariant of boundary links*, Preprint (2001), http://xxx.lanl.gov/abs/math.GT/0105028 v2.

[15] **S Garoufalidis**, **J Levine**, *Homology surgery and invariants of 3–manifolds*, Geometry and Topology Vol. 5 (2001) 75–108.

[16] **K Habiro**, *Claspers and finite type invariants of links*, Geometry and Topology Vol. 4 (2000) 1–83.

[17] **N Habegger**, **G Masbaum**, *The Kontsevich integral and Milnor's invariants*, Topology 39 (2000) 1253-1289.

[18] **M Kervaire**, **J Milnor**, *On 2–spheres in 4–manifolds*, Proc. Nat. Acad. Sci. Vol. 47 (1961) 1651–1657.

[19] **W Magnus**, **A Karass**, **D Solitar**, *Combinatorial group theory*, Dover Publications, Inc. (1976).

[20] **K Kobayashi**, *On a homotopy version of* 4*–dimensional Whitney's lemma*, Math. Sem. Notes 5 (1977) 109–116.

[21] **S Kojima**, *Milnor's* $\overline{\mu}$*–invariants, Massey products and Whitney's trick in* 4*–dimensions*, Topology and its applications 16 (1983) 43–60.

[22] **V Krushkal**, *Exponential separation in* 4*–manifolds*, Geometry and Topology, Vol. 4 (2000) 397–405.

[23] **S. Krushkal**, **P. Teichner**, *Alexander duality, Gropes and link homotopy*, Geometry and Topology Vol. 1 (1997) 51–69.

[24] **J Milnor**, *Link groups*, Annals of Math. 59 (1954) 177–195.

[25] **J Milnor**, *Isotopy of links*, Algebraic geometry and topology, Princeton Univ. Press (1957).

[26] **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.

[27] **G Masbaum**, **A Vaintrob**, *Milnor numbers, spanning trees and the Alexander-Conway polynomial*, (To appear in Advances in Mathematics) Preprint (2001) arXiv:math.GT/0111102.

[28] **R Schneiderman**, *Whitney towers and Gropes in 4–manifolds*, To appear in Trans. Amer. Math. Soc. (http://front.math.ucdavis.edu/math.GT/0310303).

[29] **R Schneiderman**, *Simple Whitney towers, half-gropes and the Arf invariant of a knot.* Preprint (2002) http://front.math.ucdavis.edu/math.GT/0310304.

[30] **R Schneiderman**, **P Teichner**, *Higher order intersection numbers of* 2*–spheres in* 4*–manifolds*, Alg. and Geom. Topology 1 (2001) 1–29.

[31] **R Schneiderman**, **P Teichner**, *Pulling apart* 2*–spheres in* 4*–manifolds*, In preparation.

[32] **R Schneiderman**, **P Teichner**, *Grope concordance of classical links*, In preparation.

[33] **P Teichner**, *Knots, von Neumann Signatures, and Grope Cobordism*, Proceedings of the International Congress of Math. Vol II: Invited Lectures (2002) 437–446.

[34] **T Wall**, *Surgery on Compact Manifolds*, London Math.Soc.Monographs 1, Academic Press 1970 or Second Edition, edited by A. Ranicki, Math. Surveys and Monographs 69, A.M.S.

[35] **H Whitney**, *The self intersections of a smooth* n*-manifold in* $2n$*-space*, Annals of Math. 45 (1944) 220–246.

[36] **M Yamasaki**, *Whitney's trick for three* 2*–dimensional homology classes of* 4*–manifolds*, Proc. Amer. Math. Soc. 75 (1979) 365–371.