

Whitney tower concordance of classical links

JAMES CONANT
 ROB SCHNEIDERMAN
 PETER TEICHNER

This paper computes *Whitney tower* filtrations of classical links. Whitney towers consist of iterated stages of Whitney disks and allow a tree-valued intersection theory, showing that the associated graded quotients of the filtration are finitely generated abelian groups. *Twisted Whitney towers* are studied and a new quadratic refinement of the intersection theory is introduced, measuring Whitney disk framing obstructions. It is shown that the filtrations are completely classified by Milnor invariants together with new *higher-order Sato-Levine* and *higher-order Arf invariants*, which are obstructions to framing a twisted Whitney tower in the 4-ball bounded by a link in the 3-sphere. Applications include computation of the *grope filtration*, and new geometric characterizations of Milnor's link invariants.

57M25; 57N10

1 Introduction

The general failure of the *Whitney move* is one reason why 4-dimensional manifolds are notoriously difficult to understand. A successful Whitney move is shown in Figure 1:




Figure 1: Left: A canceling pair of transverse intersections between two local sheets of surfaces in a 3-dimensional slice of 4-space. The translucent horizontal blue sheet appears entirely in this 3-dimensional ‘present’, and the red sheet appears as an arc which is assumed to extend into ‘past’ and ‘future’. Middle: A Whitney disk W pairing the intersections. Right: A Whitney move guided by W eliminates the intersection pair, without creating any new intersections.

In higher dimensions, this move is the key to Whitney’s strong embedding theorem [44] as well as the s -cobordism theorem and the surgery exact sequence. In each case, a pair of intersection points between two submanifolds is removed by a homotopy along an embedded Whitney disk W as in Figure 1. Whitney disks can be found by controlling the intersections between the submanifolds algebraically over the fundamental group of the ambient manifold, and in dimensions greater than four can be assumed by general position to be embedded and disjoint from the relevant submanifolds.

In four dimensions, generic intersections between Whitney disks and surface sheets can obstruct a successful Whitney move: Figure 2(a) shows how such an intersection point leads to an *unsuccessful* Whitney move.




Figure 2: (a) This Whitney move eliminates the previous pair of intersections but creates a new pair of intersections between the translucent blue sheet and the sheet described by the green arc. (b) Higher-order intersections and Whitney disks. All arcs are assumed to extend into past and future, describing local sheets of surfaces in a 4–ball.

In a 4–manifold, a Whitney disk W can always be made to be embedded (and *framed*, see section 2.2) at the cost of creating intersections with the surface sheets paired by W , but the converse is not always possible (as explained below in Section 2). It is natural to try to eliminate such *higher-order* intersections using (higher-order) Whitney moves, and this attempt leads to the notion of a *Whitney tower* (Definition 2.3), constructed on immersed surfaces in a 4–manifold by pairing up as many intersection points as possible with iterated Whitney disks (see Figure 3(a)).

A Whitney tower has a fundamental complexity called *order*, which roughly corresponds to the number of layers of Whitney disks, and in [9, 38] we introduced an accompanying obstruction theory that will be used in this paper to study link concordance. The main idea is that the chaos of multiple Whitney disks and intersection points can be organized by associating certain univalent trees to the unpaired intersections




Figure 3: (a) Part of a Whitney tower \mathcal{W} . (b) The unpaired intersections determine trivalent trees, and \mathcal{W} can be ‘split’ so that all singularities are contained in neighborhoods of embeddings of these trees.

(Figure 3), and summing these trees defines an intersection invariant whose vanishing is sufficient to raise the order of the Whitney tower by finding another layer of Whitney disks. By taking the values of this invariant in an appropriate target group, we will also be able to measure the failure of the Whitney move by determining when the order of a Whitney tower can *not* be raised.

One reason that classical links are especially relevant to Whitney towers comes from the observation that the boundaries of the embedded disk-sheets (blue, red and green) in the 4–ball described by Figure 2(a) form the Borromean rings in the boundary the 3–sphere (see e.g. Figure 15 in Section 3). Figure 4 shows how the trivalent tree associated to the unpaired intersection in the Whitney disk is preserved under the Whitney move, hinting at the fact that this tree represents an obstruction to the existence of disjointly embedded disks bounded by the Borromean rings.




Figure 4: A Whitney move preserves the associated tree.

In fact, any Whitney tower can be ‘split’ as in Figure 3(b) so that all singularities are contained in 4–ball neighborhoods of the univalent trees, each containing embedded disks bounded by links which are iterated Bing-doubles of unknots. While the local “tree-preserving” property of Figure 4 holds for a Whitney move on any Whitney disk containing an unpaired intersection, it turns out that if a Whitney disk W contains a boundary arc of a higher-order Whitney disk then a Whitney move on W locally converts the tree to two new trees, expressing a Jacobi identity (also called an IHX relation) [36, Lem.7.2]. Thus, interpreting these trees as well-defined obstructions to raising the order of a Whitney tower requires some care, see Definition 1.2.

On the other hand, if a link bounds immersed disks supporting a Whitney tower with *no* unpaired intersections, then doing Whitney moves on all the Whitney disks leads to disjointly embedded *slice disks* bounded by the link. Motivated by the notion that a Whitney tower of larger order is in some sense a “better approximation” of slice disks, the main goal of this paper is to provide an answer to the following question for any given n : “Which links in the 3–sphere bound an order n Whitney tower in the 4–ball?”

The rest of this introduction will mostly be concerned with outlining the answer to this question which is roughly summarized by the following theorem:

Theorem 1.1 *A link bounds a Whitney tower of order n if and only if its Milnor invariants, higher-order Sato-Levine invariants and higher-order Arf invariants vanish up to order n . Compare Corollary 5.11.*

We work in the smooth oriented category, with orientations frequently suppressed from notation. Throughout this paper identical statements hold in the locally flat topological category, as explained in Remark 2.1.

1.1 The Whitney tower filtration of classical links

Referring to Section 2 for a precise definition of Whitney tower, including detailed discussion of *framing* requirements, we consider the (framed) *Whitney tower filtration*

$$\cdots \subseteq \mathbb{W}_3 \subseteq \mathbb{W}_2 \subseteq \mathbb{W}_1 \subseteq \mathbb{W}_0 \subseteq \mathbb{L}$$

on the set $\mathbb{L} = \mathbb{L}(m)$ of m -component framed links in S^3 . Here $\mathbb{W}_n = \mathbb{W}_n(m)$ is the subset of those framed links that bound order n (framed) Whitney towers in B^4 . In section 1.5 we compare this filtration to other known iterated disk constructions.

Throughout this paper the number m of link components will frequently be suppressed from notation because it is fixed in most constructions.

The intersection of all \mathbb{W}_n contains all slice links because a 2-disk is a Whitney tower of arbitrarily large order. In fact, this filtration factors through link concordance; and we shall use this fact implicitly in various places.

Whitney towers built on immersed annuli connecting link components in $S^3 \times I$ induce equivalence relations of *Whitney tower concordance* on links. The quotient \mathbb{W}_n of \mathbb{W}_n modulo the equivalence relation of Whitney tower concordance of order $n + 1$ is the *associated graded* of our filtration in the sense that $L \in \mathbb{W}_{n+1}$ if and only if $L \in \mathbb{W}_n$ and $[L] = 0 \in \mathbb{W}_n$. (We will show in Section 3 that connected sum leads to a well defined group structure on \mathbb{W}_n , so 0 corresponds to the unlink.)

As a first step towards our goal of describing a classification of this filtration and the associated \mathbb{W}_n , we recall (e.g. from [9, 12, 30, 38]) a combinatorially defined group which is a natural target for the intersection invariant associated to the obstruction theory for Whitney towers.

Definition 1.2 In this paper, 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. The *order* of a tree is the number of trivalent vertices. Univalent vertices will usually be labeled from the set $\{1, 2, 3, \dots, m\}$ indexing the link components, and we consider trees up to isomorphisms preserving these labelings. Define $\mathcal{T} = \mathcal{T}(m)$ to be the free abelian group on such trees, modulo the antisymmetry (AS) and Jacobi (IHX) relations shown in Figure 5.

$$\begin{array}{l} \text{AS:} \\ \text{IHX:} \end{array} \quad \begin{array}{c} \begin{array}{c} \diagup \quad \diagdown \\ \quad \quad \quad | \\ \quad \quad \quad \diagdown \quad \diagup \end{array} + \begin{array}{c} \diagup \quad \diagdown \\ \quad \quad \quad | \\ \quad \quad \quad \diagdown \quad \diagup \\ \quad \quad \quad \diagdown \quad \diagup \end{array} = 0 \\ \\ \begin{array}{c} \diagup \quad \diagdown \\ \quad \quad \quad | \\ \quad \quad \quad \diagdown \quad \diagup \end{array} - \begin{array}{c} \diagup \quad \diagdown \\ \quad \quad \quad | \\ \quad \quad \quad \diagdown \quad \diagup \end{array} + \begin{array}{c} \diagup \quad \diagdown \\ \quad \quad \quad | \\ \quad \quad \quad \diagdown \quad \diagup \\ \quad \quad \quad \diagdown \quad \diagup \end{array} = 0 \end{array}$$

Figure 5: Local pictures of the *antisymmetry* (AS) and *Jacobi* (IHX) relations in \mathcal{T} . Here all trivalent orientations are induced from a fixed orientation of the plane, and univalent vertices possibly extend to subtrees which are fixed in each equation.

Since the AS and IHX relations are homogeneous with respect to order, \mathcal{T} inherits a grading $\mathcal{T} = \bigoplus_n \mathcal{T}_n$, where $\mathcal{T}_n = \mathcal{T}_n(m)$ is the free abelian group on order n trees, modulo AS and IHX relations.

As recalled below in Section 2, the Whitney tower obstruction theory [9, 38] assigns to each order n Whitney tower \mathcal{W} an order n intersection invariant $\tau_n(\mathcal{W}) \in \mathcal{T}_n$, which is defined by summing the trees pictured in Figure 3(b) (see Figure 6 for an explicit example in the 4–ball). The tree orientations are induced by Whitney disk orientations via a convention that corresponds to the AS relations (section 2.3), and the IHX relations can be realized geometrically by controlled maneuvers on Whitney towers as described in [9, 35]. It follows from the obstruction theory that a link bounds a Whitney tower \mathcal{W} of order n with $\tau_n(\mathcal{W}) = 0$ if and only if it bounds a Whitney tower of order $n + 1$. This fact is the essential ingredient in the proof of the following theorem (see Section 3):

Theorem 1.3 *The sets W_n are finitely generated abelian groups under the (well-defined) operation of connected sum $\#$ and there are epimorphisms $R_n : \mathcal{T}_n \twoheadrightarrow W_n$.*

These *realization maps* R_n are defined similarly to Cochran’s iterated Bing-doubling construction for realizing Milnor invariants [2], and are equivalent to ‘simple clasper surgery along trees’ in the sense of Goussarov and Habiro (see section 3.3, and Figure 6 for an example).




Figure 6: The realization map R_2 sends the tree t on the left to the link $L \subset S^3$ shown in the middle. The trace of a null-homotopy of L described by a pair of crossing-changes between the blue component 1 and the red component 2 supports an order 2 Whitney tower $\mathcal{W} \subset B^4$ bounded by L , with $\tau_2(\mathcal{W}) = t$, as shown on the right. (Pushing further into B^4 would show a 3-component unlink bounding disjointly embedded disks).

Theorem 4 of [38] implies that R_n is *rationally* an isomorphism for all n , and from this a formula for the free ranks of the groups W_n can be given. Until now we didn’t know much about torsion phenomena, and the following fact comes as a huge surprise:

Theorem 1.4 *In all even orders, the realization maps $R_{2k} : \mathcal{T}_{2k} \rightarrow W_{2k}$ are isomorphisms and W_{2k} are free abelian groups of known rank, detected by Milnor invariants.*

As explained in section 5 (Theorem 5.5), this result follows from the relationship between the Whitney tower intersection invariant τ and Milnor invariants [11, 38], together with our proof in [12] of a combinatorial conjecture of J. Levine [29], which also implies that \mathcal{T}_{2k} is a free abelian group of known rank [12, Thm.1.5].

The affirmation of Levine's conjecture also implies that the torsion in \mathcal{T}_{2k-1} is generated by symmetric trees of the form $i \prec J$, where J is a subtree of order $k-1$, and i is a univalent vertex label (see [12, Cor.1.2]). These trees are actually 2-torsion by the antisymmetry relation and hence all torsion in \mathcal{T} is 2-torsion. The next result shows that a large part of this 2-torsion actually maps trivially to W_{2k-1} .

Theorem 1.5 *The realization maps R_{2k-1} factor through a quotient $\tilde{\mathcal{T}}_{2k-1}$ of \mathcal{T}_{2k-1} .*

Theorem 1.5 is proved in Section 4.4, where it is also shown that our Whitney tower obstruction theory descends to these *reduced groups* $\tilde{\mathcal{T}}_{2k-1}$:

Definition 1.6 Let $\tilde{\mathcal{T}}_{2k-1} := \mathcal{T}_{2k-1} / \text{Im } \Delta_{2k-1}$, where $\Delta_{2k-1}: \mathcal{T}_{k-1} \rightarrow \mathcal{T}_{2k-1}$ is defined on generators t of order $k-1$ as follows. For any univalent vertex v of t , denote by $\ell(v)$ the label of v , and write $t = \ell(v) \prec T_v(t)$. Then we get a 2-torsion element of \mathcal{T}_{2k-1} defined by

$$\Delta_{2k-1}(t) := \sum_v \ell(v) \prec \begin{matrix} T_v(t) \\ T_v(t) \end{matrix}$$

where the sum is over all univalent vertices v of t .

Figure 7: The framing relations in orders 1 and 3.

We refer to the relations $\Delta_{2k-1}(t) = 0$ as *framing relations* because they correspond to the image of *twisted IHX* relations in a *twisted Whitney tower* via a conversion to a framed Whitney tower, as explained in Section 4.4.

Conjecturally, all odd order *reduced realization maps* $\tilde{R}_{2k-1}: \tilde{\mathcal{T}}_{2k-1} \rightarrow W_{2k-1}$ are isomorphisms, and the following theorem confirms this in half of the cases:

Theorem 1.7 *The reduced realization maps \tilde{R}_{4k-1} are isomorphisms and the torsion of W_{4k-1} is a \mathbb{Z}_2 -vector space of known dimension, detected by higher order Sato-Levine invariants.*

The *higher-order Sato-Levine invariants* are certain projections of Milnor invariants, shifted down one order. They represent obstructions to framing a twisted Whitney tower, as explained in section 5. To outline the proof of Theorem 1.7, we next introduce the *twisted Whitney tower filtration*, and explain how *higher-order Arf invariants* play a role in completing the classifications of both the twisted and framed filtrations. Theorem 1.7 will follow from Theorem 5.9 in section 5.

1.2 The twisted Whitney tower filtration of classical links

As illustrated in Figure 1 and detailed below in section 2.2, a successful Whitney move requires the existence of disjoint parallel copies of the Whitney disk (which extend a canonical normal section over the boundary). Such a Whitney disk is said to be *framed*, and this condition is required for all Whitney disks in a Whitney tower. An order n *twisted Whitney tower* is the same as an order n (framed) Whitney tower, except that the framing condition is not required for Whitney disks of order greater than or equal to $n/2$.

Denote by $\mathbb{W}_n^\infty = \mathbb{W}_n^\infty(m)$ the set of framed m -component links that bound order n *twisted Whitney towers*. This gives the *twisted Whitney tower filtration*

$$\cdots \subseteq \mathbb{W}_3^\infty \subseteq \mathbb{W}_2^\infty \subseteq \mathbb{W}_1^\infty \subseteq \mathbb{W}_0^\infty = \mathbb{L}$$

We refer to Section 3 for a precise definition, also of the associated graded $W_n^\infty = \mathbb{W}_n^\infty(m)$; and to Section 2 for details on twisted Whitney towers, including the associated *twisted intersection invariant* $\tau_n^\infty(\mathcal{W}) \in \mathcal{T}_n^\infty$ in section 2.4. These notions appear in this paper for the first time.

Briefly, the odd order groups $\mathcal{T}_{2k-1}^\infty$ are defined as quotients of \mathcal{T}_{2k-1} by the torsion subgroups, generated by trees of the form $i \prec J$; where J is a subtree of order $k-1$, and i is a univalent vertex label. These *boundary-twist relations* correspond to the intersections created by performing a boundary-twist on an order k Whitney disk (Figure 18).

In addition to the generating trees for \mathcal{T}_{2k} , the groups \mathcal{T}_{2k}^∞ include trees of the form

$$\infty \text{ --- } J$$

with k trivalent vertices and one univalent vertex labeled by the twist symbol ∞ (whereas all other univalent vertices are still labeled by elements in $\{1, \dots, m\}$). These additional generators are called ∞ -trees or *twisted trees* because the vertex labeled by ∞ will lie in a twisted Whitney disk. They represent framing obstructions on order k

Whitney disks and are involved in the new *symmetry*, *twisted IHX*, and *interior twist* relations, all of which have geometric interpretations (Definition 2.8).

Note that here and in the following the symbol ∞ represents a *twist*, and in particular does *not* stand for “infinity”.

Theorem 1.8 *The sets W_n^∞ are finitely generated abelian groups under the (well-defined) operation of connected sum $\#$ and there are epimorphisms $R_n^\infty : \mathcal{T}_n^\infty \twoheadrightarrow W_n^\infty$.*

These *twisted realization maps* are constructed in Section 3. As in the framed setting, the key to the proof is a criterion for *raising the order of a twisted Whitney tower*:

Theorem 1.9 *A link bounds a twisted Whitney tower \mathcal{W} of order n with $\tau_n^\infty(\mathcal{W}) = 0$ if and only if it bounds a twisted Whitney tower of order $n + 1$.*

Theorem 1.9 follows from the more general Theorem 2.10, which is proved in Section 4 using *twisted Whitney moves* (Lemma 4.1) as well as boundary-twists and a construction for geometrically canceling twisted Whitney disks.

Next we bring Milnor’s μ -invariants into the picture. Let $L_n = L_n(m)$ be the degree n part of the free Lie \mathbb{Z} -algebra $L = \bigoplus L_n$ generated by degree 1 generators $\{X_1, \dots, X_m\}$. Via the usual correspondence between rooted oriented labeled trees and non-associative bracketings, L_n can be identified with the abelian group on *rooted* trees of order $(n - 1)$, where the root is a chosen unlabeled univalent vertex and the other univalent vertices are labeled as before from the index set $\{1, 2, 3, \dots, m\}$, modulo IHX relations (Jacobi identities) and *self-annihilation relations* which kill all generators having an order two symmetry. The self-annihilation relations, which are expressed in terms of brackets as $[X, X] = 0$, imply the antisymmetry relations of Figure 5, but not vice versa.

Definition 1.10 Define $\eta_n : \mathcal{T}_n^\infty \rightarrow L_1 \otimes L_{n+1}$ on trees t by $\eta_n(t) := \sum_{v \in t} X_{\ell(v)} \otimes T_v(t)$, where the sum is over all univalent vertices v of t , with $T_v(t)$ denoting the rooted tree gotten by replacing v with a root, and $\ell(v)$ the original label of v . On ∞ -trees $\infty - J$, define $\eta_n(\infty - J) := \frac{1}{2}\eta_n(J - J)$, where dividing by 2 makes sense since the terms in $\eta_n(J - J)$ all have even coefficients. It is easy to check that η_n is a well-defined homomorphism for all n (see [11, 13]).

Define the group D_n to be the kernel of the bracket map $L_1 \otimes L_{n+1} \rightarrow L_{n+2}$ which sends $X_i \otimes Y \mapsto [X_i, Y]$. This group D_n is the natural target for the first non-vanishing Milnor invariants of a link [23, 33], and it turns out that η_n maps \mathcal{T}_n^∞ onto D_n [11]. In [11] we explain the precise relationship between twisted Whitney towers and the (first non-vanishing) Milnor invariants via a 4-dimensional incarnation of the η_n -map using *grope duality* [27]. In particular, we have the following result:

Theorem 1.11 ([11]) *The Milnor invariants of length $\leq n + 1$ vanish for links $L \in \mathbb{W}_n^\infty$, and the length $n + 2$ Milnor invariants of L gives rise to a homomorphism $\mu_n: \mathbb{W}_n^\infty \rightarrow \mathbb{D}_n$ such that $\mu_n \circ R_n^\infty = \eta_n$.*

We will refer to the maps $\mu_n: \mathbb{W}_n^\infty \rightarrow \mathbb{D}_n$ as *order n Milnor invariants*. For $n = 0$ they give linking numbers as well as framings (as coefficients of $X_i \otimes X_i$).

Corollary 1.12 *There is a commutative diagram of epimorphisms*

$$\begin{array}{ccc} \mathcal{T}_n^\infty & \xrightarrow{R_n^\infty} & \mathbb{W}_n^\infty \\ & \searrow \eta_n & \downarrow \mu_n \\ & & \mathbb{D}_n \end{array}$$

Moreover, we show below in Theorem 6.5 that $\eta_n: \mathcal{T}_n^\infty \rightarrow \mathbb{D}_n$ is an *isomorphism* except when $n \equiv 2 \pmod{4}$. Since \mathbb{D}_n is a free abelian group of known rank for all n [33], this completes the computation of \mathbb{W}_n^∞ in three quarters of the cases:

Theorem 1.13 *If $n \not\equiv 2 \pmod{4}$, the maps R_n^∞ and μ_n give rise to isomorphisms*

$$\mathcal{T}_n^\infty \cong \mathbb{W}_n^\infty \cong \mathbb{D}_n$$

The last quarter of the cases is more complicated as can already be seen for $n = 2$: In the case $m = 1$ of knots, we show in [11] that the classical Arf invariant induces an isomorphism $\mathbb{W}_2^\infty(1) \cong \mathbb{Z}_2$, whereas all Milnor invariants vanish for knots.

In Corollary 6.6 of section 6.3 we derive the following result which computes the kernel of η_n for all $n \equiv 2 \pmod{4}$:

Proposition 1.14 *The map sending $1 \otimes J$ to $\infty \prec_J^J \in \mathcal{T}_{4k-2}^\infty$ for rooted trees J of order $k - 1$ defines an isomorphism $\mathbb{Z}_2 \otimes \mathbb{L}_k \cong \text{Ker}(\eta_{4k-2}: \mathcal{T}_{4k-2}^\infty \rightarrow \mathbb{D}_{4k-2})$.*

It follows from Corollary 1.12 that $\mathbb{Z}_2 \otimes \mathbb{L}_k$ is also an upper bound on the kernels of the epimorphisms R_{4k-2}^∞ and μ_{4k-2} , and the calculation of \mathbb{W}_{4k-2}^∞ will be completed by invariants defined on the kernel of μ_{4k-2} which we believe are new concordance invariants generalizing the classical Arf invariant, as we describe next.

Higher-order Arf invariants

Let K_{4k-2}^∞ denote the kernel of $\mu_{4k-2} : W_{4k-2}^\infty \rightarrow D_{4k-2}$. It follows from Corollary 1.12 and Proposition 1.14 above that mapping $1 \otimes J$ to $R_{4k-2}^\infty(\infty \prec J)$ induces a surjection $\alpha_k^\infty : \mathbb{Z}_2 \otimes L_k \rightarrow K_{4k-2}^\infty$, for all $k \geq 1$. Denote by $\overline{\alpha}_k^\infty$ the induced isomorphism on $(\mathbb{Z}_2 \otimes L_k) / \text{Ker } \alpha_k^\infty$.

Definition 1.15 The *higher-order Arf invariants* are defined by

$$\text{Arf}_k := (\overline{\alpha}_k^\infty)^{-1} : K_{4k-2}^\infty \rightarrow (\mathbb{Z}_2 \otimes L_k) / \text{Ker } \alpha_k^\infty$$

From Corollary 1.12, Theorem 1.13, Proposition 1.14 and Definition 1.15 we see that the groups W_n^∞ are computed by the Milnor and higher-order Arf invariants:

Corollary 1.16 The groups W_n^∞ are classified by Milnor invariants μ_n and, in addition, higher-order Arf invariants Arf_k for $n = 4k - 2$.

In particular, it follows that a link bounds an order n twisted Whitney tower if and only if its Milnor invariants and higher-order Arf invariants vanish up to order n .

We conjecture that the α_k^∞ are isomorphisms, which would mean that the Arf_k are very interesting new concordance invariants:

Conjecture 1.17 $\text{Arf}_k : K_{4k-2}^\infty \rightarrow \mathbb{Z}_2 \otimes L_k$ are isomorphisms for all k .

Conjecture 1.17 would imply that $W_{4k-2}^\infty \cong \mathcal{T}_{4k-2}^\infty \cong (\mathbb{Z}_2 \otimes L_k) \oplus D_{4k-2}$ where the second isomorphism (is non-canonical and) already follows from Proposition 1.14. Conjecture 1.17 is true for $k = 1$, with Arf_1 given by the classical Arf invariants of the link components [11]. It remains an open problem whether Arf_k is non-trivial for any $k > 1$. The links $R_{4k-2}^\infty(\infty \prec J)$ 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]. Such links are known not to be slice by work of J.C. Cha [1], providing evidence in support of Conjecture 1.17.

In combination with Theorem 1.13, Conjecture 1.17 can be succinctly expressed in terms of the twisted Whitney tower filtration classification as the statement: “the twisted realization maps $R_n^\infty : \mathcal{T}_n^\infty \rightarrow W_n^\infty$ are isomorphisms for all n .”

A table of the groups $W_n^\infty(m)$ for low values of n, m is given in Figure 8, where the higher-order Arf invariant Arf_2 appears in order 6. The currently unknown ranks of

	1	2	3	4	5
0	\mathbb{Z}	\mathbb{Z}^3	\mathbb{Z}^6	\mathbb{Z}^{10}	\mathbb{Z}^{15}
1	0	0	\mathbb{Z}	\mathbb{Z}^4	\mathbb{Z}^{10}
2	\mathbb{Z}_2	$\mathbb{Z} \oplus \mathbb{Z}_2^2$	$\mathbb{Z}^6 \oplus \mathbb{Z}_2^3$	$\mathbb{Z}^{20} \oplus \mathbb{Z}_2^4$	$\mathbb{Z}^{50} \oplus \mathbb{Z}_2^5$
3	0	0	\mathbb{Z}^6	\mathbb{Z}^{36}	\mathbb{Z}^{126}
4	0	\mathbb{Z}^3	\mathbb{Z}^{28}	\mathbb{Z}^{146}	\mathbb{Z}^{540}
5	0	0	\mathbb{Z}^{36}	\mathbb{Z}^{340}	\mathbb{Z}^{1740}
6	0	$\mathbb{Z}^6 \oplus \mathbb{Z}_2^{e_2}$	$\mathbb{Z}^{126} \oplus \mathbb{Z}_2^{e_3}$	$\mathbb{Z}^{1200} \oplus \mathbb{Z}_2^{e_4}$	$\mathbb{Z}^{7050} \oplus \mathbb{Z}_2^{e_5}$

Figure 8: A table of the groups $W_n^\infty(m)$, where m runs horizontally and n runs vertically. The possible ranges of the torsion exponents in order 6 depend on the currently unknown ranks of Arf_2 : $0 \leq e_2 \leq 1$, $0 \leq e_3 \leq 3$, $0 \leq e_4 \leq 6$, $0 \leq e_5 \leq 10$.

Arf_2 are represented by the ranges of possible ranks of the 2-torsion subgroups of the groups $W_6^\infty(m)$.

For $n = 0$, the groups are freely generated by the image under R_0^∞ of trees $i - j$, with $i \neq j$, and twisted trees $\infty - j$. The resulting links are detected by linking numbers and framings, respectively. For order $n = 1$, the generators come (via R_1^∞) from trees $i \leftarrow_k^j$ where all indices are distinct (otherwise the tree is zero in \mathcal{T}_1^∞ by the boundary-twist relations). They are detected by Milnor's triple invariants $\mu(ijk)$.

In order $n = 2$, generators include (R_2^∞ of) ∞ -trees $\infty \leftarrow_k^i$ (recall that these indeed lie in \mathcal{T}_2^∞ even though the tree has only one trivalent vertex). If $i \neq j$ these are of infinite order, detected by Milnor's $\mu(ijij)$, but for $i = j$ they have order 2 and are detected by the classical Arf invariant of the i th component. This shows how our groups $\mathcal{T}_{4k-2}^\infty$ combine Milnor and Arf invariants in one new formalism.

1.3 Framing twisted Whitney towers

Translation of the classification of the twisted Whitney tower filtration back into the framed setting will be accomplished in Section 5 using a new interpretation of certain first non-vanishing Milnor invariants as obstructions to framing a twisted Whitney tower. These are the higher-order Sato-Levine invariants which are defined in all odd orders of the framed Whitney tower filtration (section 5.1). The higher-order Arf invariants will also appear as framing obstructions (section 5.2), however they will be shifted down one order, due to the fact that a twisted Whitney tower of order $2k$ can always be converted into a framed Whitney tower of order $2k - 1$ by twisting and IHX

constructions (section 4.3). These geometric constructions will explain the origin of the *framing relations* introduced above in Definition 1.6.

Setting $\tilde{\mathcal{T}}_{2k} := \mathcal{T}_{2k}$ in even orders, Theorem 5.9 will show that the reduced realization maps $\tilde{R}_n : \tilde{\mathcal{T}}_n \rightarrow W_n$ are isomorphisms in three quarters of the cases, in close analogy with Theorem 1.13 above. Then the higher-order Arf invariants will again appear in the other quarter of cases, and Conjecture 1.17 will have an analogous expression in terms of the framed Whitney tower filtration classification as the statement: “the realization maps $\tilde{R}_n : \tilde{\mathcal{T}}_n \rightarrow W_n$ are isomorphisms for all n ” (section 5.3).

However, the analogy with Theorem 1.13 does *not* hold for the Milnor invariants μ_n in the framed filtration, leading to the appearance of the higher-order Sato-Levine invariants in the classification of the framed filtration described in Corollary 5.11. This subtle interaction between Milnor invariants and framing obstructions is the reason why the framed classification is trickier to describe.

A table of the framed filtration groups $W_n(m)$ for low values of n, m is given in Figure 9, where the higher-order Arf invariant Arf_2 appears in order 5. The higher-order Sato-Levine invariants correspond to 2-torsion in all odd orders (for $m > 1$), and the ranges of possible ranks of the 2-torsion subgroups of the groups $W_5(m)$ correspond to the possible ranks of Arf_2 (as in Figure 8).

	1	2	3	4	5
0	\mathbb{Z}	\mathbb{Z}^3	\mathbb{Z}^6	\mathbb{Z}^{10}	\mathbb{Z}^{15}
1	\mathbb{Z}_2	\mathbb{Z}_2^3	$\mathbb{Z} \oplus \mathbb{Z}_2^6$	$\mathbb{Z}^4 \oplus \mathbb{Z}_2^{10}$	$\mathbb{Z}^{10} \oplus \mathbb{Z}_2^{15}$
2	0	\mathbb{Z}	\mathbb{Z}^6	\mathbb{Z}^{20}	\mathbb{Z}^{50}
3	0	\mathbb{Z}_2^2	$\mathbb{Z}^6 \oplus \mathbb{Z}_2^8$	$\mathbb{Z}^{36} \oplus \mathbb{Z}_2^{20}$	$\mathbb{Z}^{126} \oplus \mathbb{Z}_2^{40}$
4	0	\mathbb{Z}^3	\mathbb{Z}^{28}	\mathbb{Z}^{146}	\mathbb{Z}^{540}
5	0	$\mathbb{Z}_2^{e_2}$	$\mathbb{Z}^{36} \oplus \mathbb{Z}_2^{e_3}$	$\mathbb{Z}^{340} \oplus \mathbb{Z}_2^{e_4}$	$\mathbb{Z}^{1740} \oplus \mathbb{Z}_2^{e_5}$
6	0	\mathbb{Z}^6	\mathbb{Z}^{126}	\mathbb{Z}^{1200}	\mathbb{Z}^{7050}

Figure 9: A table of the groups $W_n(m)$, where m runs horizontally and n runs vertically. The possible ranges of the torsion exponents in order 5 depend on the currently unknown ranks of Arf_2 : $3 \leq e_2 \leq 4$, $18 \leq e_3 \leq 21$, $60 \leq e_4 \leq 66$, $150 \leq e_5 \leq 160$.

For $n = 0$, the groups come from trees $i - j$, and are detected by linking numbers for $i \neq j$ and framings for $i = j$. For order $n = 1$, the generators come (via R_1) from trees $i \leftarrow_k^j$. If all indices are distinct then they are detected by Milnor’s triple invariants $\mu(ijk)$. However, in $\tilde{\mathcal{T}}_1$ repeating indices also give nontrivial elements of order 2. If

$i = j = k$, these are detected by the classical Arf invariant of the i th component. In the case where exactly two indices are equal, one needs the classical Sato-Levine invariant (but has to note the framing relations from Figure 7).

The main tool for deriving the framed classification from the twisted one is a commutative diagram of exact sequences (Theorem 5.1) in which the various realization maps connect the tree-groups to the associated graded groups of the filtrations. As a consequence of our resolution in [12] of the Levine Conjecture (Theorem 5.4), all the relevant tree-groups are completely computed. So together with some additional geometric and algebraic arguments, the graded groups associated to the framed filtration can be computed in terms of those of the twisted filtration.

In Section 6, the diagram of Theorem 5.1 relating the \mathcal{T} - and W -groups is extended by the relevant η - and μ -maps to include exact sequences of D -groups, giving a bird's eye view of the classifications. The resulting pair of *master diagrams* gives a succinct summary of the overall algebraic structure connecting the \mathcal{T} -, W -, and D -groups.

1.4 Applications to gropes and k -slice theorems

Recall (e.g. from [42]) that a *grope of class k* is defined recursively as follows: A grope of class 1 is a circle and a grope of class 2 is an orientable surface Σ with one boundary component. A grope of class k is formed by attaching to every dual pair of curves in a symplectic basis for Σ a pair of gropes whose classes add to k . For details on gropes, including framing conditions, see e.g. [7, 8, 9, 19, 20, 26, 27, 34, 41].

The grope filtration by class

The *grope filtration* (by class) on the set $\mathbb{L} = \mathbb{L}(m)$ of framed links in S^3 with m components is defined by the sets $\mathbb{G}_n \subset \mathbb{L}$ of framed links whose components bound class $(n+1)$ disjointly embedded framed gropes in B^4 . The index shift is explained by the next theorem, and for the same reason we *define* \mathbb{G}_0 to be the set of evenly framed links. The main result from [34] implies that this grope filtration equals our framed Whitney tower filtration and hence is being computed in this paper:

Theorem 1.18 ([34]) $\mathbb{G}_n = \mathbb{W}_n$ for all n .

The following result is a sample geometric application of our computations, characterizing links with certain vanishing Milnor invariants. Details and related applications are described in [11].

Theorem 1.19 ([11]) *A link has vanishing Milnor invariants of all orders $\leq 2k - 2$ (lengths $\leq 2k$) if and only if its components bound disjointly embedded surfaces Σ_i in the 4–ball, with each surface a connected sum of two surfaces Σ'_i and Σ''_i such that*

- (i) *a symplectic basis of curves on Σ'_i bound disjointly embedded framed gropes $G_{i,j}$ of class k in the complement of $\Sigma := \cup_i \Sigma_i$, and*
- (ii) *a symplectic basis of curves on Σ''_i bound immersed disks in the complement of $\Sigma \cup G$, where G is the union of all $G_{i,j}$.*

Theorem 1.19 is a considerable strengthening of the Igusa-Orr *k-slice Theorem* [25]: Since the geometric conditions in both theorems are equivalent to the vanishing of Milnor’s invariants through order $2k - 2$ (length $2k$), one can read our result as saying that the *immersed gropes* of class k found by Igusa and Orr can be cleaned up to immersed *disks* (these are immersed gropes of arbitrarily high class) or *embedded* gropes of class k .

1.5 Comparisons with other iterated disk constructions

Andrew Casson was the first who tried to recover the Whitney move in dimension four by an iterated disk construction. He started with a simply connected 4–manifold M with a knot K in its boundary. He looked for conditions so that K would bound an embedded disk in M . His starting point was an *algebraically transverse sphere* for a (singular) disk in M bounding K , an assumption that is satisfied in the setting of the s-cobordism theorem or the surgery exact sequence (but not for $M = B^4$). He then showed that K bounds a *Casson tower* of arbitrary height in M . In such a tower, one attaches an immersed disk to the accessory circles of every intersection point in a previous stage (and requires that the new disk does not intersect previous stages).

Mike Freedman [16] realized that one can actually re-embed one Casson tower into another and that one can obtain enough geometric control to prove his breakthrough result: Any Casson tower of height > 3 contains in its neighborhood a topologically-flat embedded disk with boundary K . This implies Freedman’s classification result for simply connected closed 4–manifolds and leads to many stunning applications.

However, there can be no obstruction theory for finding Casson towers of larger and larger height, not even in $M = B^4$ (where a transverse sphere cannot exist): Any knot K bounds a Casson tower of height 1 (which is just a singular disk) and if K bounds a Casson tower of height 4 then it is topologically slice (and hence bounds a Casson tower of arbitrary height).

This motivated Cochran, Orr and the third author [5] to study another type of tower, now best called a *symmetric* Whitney tower of *height* n . Here one inductively attaches Whitney disks to previous stages but only allows these new Whitney disks to intersect each other (and not the previously constructed stages). It follows that a symmetric Whitney tower of height h is a (particularly nice) Whitney tower of order 2^h as studied in this paper, see [34].

Such symmetric Whitney towers have an extremely rich theory, even in the case of knots (see [6] for the fact that the filtration is nontrivial for all heights). All the iterated graded groups are in fact infinitely generated [4], one reason being the existence of higher-order von Neumann signatures that take values in the reals \mathbb{R} (infinitely generated as abelian group). There are currently no known algebraic criteria for raising the height of a symmetric Whitney tower, and hence not too much hope for a complete classification of the symmetric Whitney tower filtration of links, or even knots.

That's why the current authors set out to study a simpler version of this filtration, and succeeded in giving the first instance of a complete computation of a filtration defined via an iterated disk construction, as described in this paper. These Whitney tower filtrations have analogues for immersed 2-spheres in 4-manifolds, including a formulation of the proposed higher-order Arf invariants. The order 1 theory goes back to [18] (see also [31, 37, 40, 43], and 10.8A and 10.8B of [19] where the relation to the Kirby-Siebenmann invariant is explained), but the higher-order theory is not generally understood for closed 4-manifolds.

The relationship between Milnor invariants and trees goes back to Cochran's method of constructing links realizing given (integer) Milnor invariants by "Bing-doubling along a tree" [2, 3]. It is intriguing to note that the sequences of circles of intersection between Seifert surfaces that arise in Cochran's construction are strongly suggestive of the Whitney disk boundaries that appear in our construction in section 3.3.

The classifications described here also have implications for 3-manifolds, in particular filtrations on homology cylinders, as described in [14], which also explores the Whitney tower filtrations on string links and their relation to finite type invariants.

In a future paper we will show that the relation of order n Whitney tower concordance on links is the same as the equivalence relation generated by concordance and simple clasper surgeries with n nodes. This is the same as an equivalence relation on string links considered by Meilhan and Yasuhara [32] called C_{n+1} -concordance. They give a list of classifying invariants up to order 4 which consists of classical Arf invariants, Milnor invariants and various mod 2 reductions of Milnor invariants (what we here call

higher-order Sato-Levine invariants). They stop just short of order 5 which is where the first higher-order Arf invariant Arf_2 lives (Figure 9).

Acknowledgments: This paper was partially written while the first two authors were visiting the third author at the Max-Planck-Institut für Mathematik in Bonn. They all thank MPIM for its stimulating research environment and generous support. The exposition of this paper was significantly improved by a careful and insightful anonymous referee. The first author was also supported by NSF grant DMS-0604351 and the last author was also supported by NSF grants DMS-0806052 and DMS-0757312. The second author was partially supported by PSC-CUNY research grant PSCREG-41-386.

Contents

1	Introduction	1
1.1	The Whitney tower filtration of classical links	4
1.2	The twisted Whitney tower filtration of classical links	8
1.3	Framing twisted Whitney towers	12
1.4	Applications to gropes and k -slice theorems	14
1.5	Comparisons with other iterated disk constructions	15
2	Whitney towers	18
2.1	Whitney disks and higher-order intersections	20
2.2	Twisted and framed Whitney disks	21
2.3	Twisted Whitney towers and their orientations	22
2.4	Intersection invariants for (twisted) Whitney towers	23
2.5	Splitting twisted Whitney towers	28
3	The realization maps	29
3.1	Band sums of links	31
3.2	Definition of the realization maps	32
3.3	Surjectivity of the realization maps	35

4	Implications of the twisted IHX construction	36
4.1	Proof of the twisted order-raising Theorem 2.10	38
4.2	The geometric twisted IHX relation	44
4.3	Twisted even order and framed odd order Whitney towers	46
4.4	Obstruction theory for the reduced tree groups	46
5	From the twisted to the framed classification	49
5.1	Higher-order Sato-Levine invariants	51
5.2	Higher-order Arf invariants in the framed filtration	53
5.3	The Arf invariant conjecture in terms of reduced realization maps . . .	53
5.4	Exactness of the tree sequence	54
6	Summary of computations for the Whitney filtrations	56
6.1	The easy Master Diagram	56
6.2	The hard Master Diagram	57
6.3	Variations on the Levine conjecture	59
	Bibliography	61

2 Whitney towers

We sketch here the relevant theory of Whitney towers as developed in [9, 34, 38], giving details for the new notion of *twisted* Whitney towers. We work in the *smooth oriented* category (with orientations usually suppressed from notation), even though all our results hold in the locally flat topological category by the basic results on topological immersions in Freedman–Quinn [19]. In fact, it can be shown that the filtrations \mathbb{W}_n , \mathbb{W}_n^∞ and \mathbb{C}_n are identical in the smooth and locally flat settings. This is because a topologically flat surface can be promoted to a smooth surface at the cost of only creating unpaired intersections of arbitrarily high order (see Remark 2.1).

Operations on trees

To describe Whitney towers it is convenient to use the bijective correspondence between formal non-associative bracketings of elements from the index set $\{1, 2, 3, \dots, m\}$ and rooted trees, trivalent and oriented, with each univalent vertex labeled by an element from the index set, except for the *root* univalent vertex which is left unlabeled. Recall from Definition 1.2 that an orientation of a tree is determined by cyclic orderings of the adjacent edges around each trivalent vertex.

Definition 2.1 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 10 upper right). The orientation of (I, J) is inherited from those of I and J as well as the order in which they are glued.
- (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 10 lower right.)
- (iii) The *order* of a tree, rooted or unrooted, is defined to be the number of trivalent vertices.

The notation of this paper will not distinguish between a bracketing and its corresponding rooted tree (as opposed to the notation I and $t(I)$ used in [34, 38]). In [34, 38] the inner product is written as a dot-product, and the rooted product is denoted by $*$.




Figure 10: 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 clockwise orientation of the plane.

2.1 Whitney disks and higher-order intersections

A collection $A_1, \dots, A_m \looparrowright (M, \partial M)$ of connected surfaces in a 4-manifold M is a *Whitney tower of order zero* if the A_i are *properly immersed* in the sense that the boundary is embedded in ∂M and the interior is generically immersed in $M \setminus \partial M$.

To each *order zero surface* A_i is associated the order zero rooted tree consisting of an edge with one vertex labeled by i , and to each transverse intersection $p \in A_i \cap A_j$ is associated the order zero tree $t_p := \langle i, j \rangle$ 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 order 1 rooted Y-tree (i, j) , 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 M , 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. (The cyclic orientation at the trivalent vertex of the bracket (i, j) corresponds to an orientation of $W_{(i,j)}$ via a convention described below in 2.3.)

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 11); with the understanding that if, say, I is just a singleton i , then W_I denotes the order zero surface A_i . Note that a Whitney disk $W_{(I,J)}$ can be created by a finger move pushing W_J through W_I .

To any transverse intersection $p \in W_{(I,J)} \cap W_K$ between $W_{(I,J)}$ and any W_K is associated the un-rooted tree $t_p := \langle (I, J), K \rangle$ (see right-hand side of Figure 11).




Figure 11: On the left, (part of) the rooted tree (I, J) associated to a Whitney disk $W_{(I,J)}$. On the right, (part of) the unrooted tree $t_p = \langle (I, J), K \rangle$ associated to an intersection $p \in W_{(I,J)} \cap W_K$. Note that p corresponds to where the roots of (I, J) and K are identified to a (non-vertex) point in $\langle (I, J), K \rangle$.

Definition 2.2 The *order of a Whitney disk* W_I is defined to be the order of the rooted tree I , and the *order of a transverse intersection* p is defined to be the order of the tree t_p .

Definition 2.3 A collection \mathcal{W} of properly immersed surfaces together with higher-order Whitney disks is an *order n Whitney tower* if \mathcal{W} contains no unpaired intersections of order less than n .

The Whitney disks in \mathcal{W} must have disjointly embedded boundaries, and generically immersed interiors. All Whitney disks and order zero surfaces must also be *framed* (as discussed in the next subsection).

Remark 2.1 We sketch here a brief explanation of why the smooth and locally flat filtrations are equal. A locally flat surface can be made smooth by a small perturbation, which after introducing cusps as necessary can be assumed to be a regular (locally flat) homotopy. By a general position argument, this regular homotopy can be assumed to be a finite number of finger moves, which are guided by arcs and lead to canceling self-intersection pairs which admit small disjointly embedded Whitney disks (which are ‘inverses’ to the finger moves). These Whitney disks are only locally flat, but can be perturbed to be smooth, again only at the cost of creating paired self-intersections, and iteration of this process leads to an arbitrarily high-order smooth sub-Whitney tower pairing all intersections created by the original surface perturbation.

2.2 Twisted and framed Whitney disks

The normal disk-bundle of a Whitney disk W in M is isomorphic to $D^2 \times D^2$, and comes equipped with a canonical nowhere-vanishing *Whitney section* over the boundary given by pushing ∂W tangentially along one sheet and normally along the other, avoiding the tangential direction of W (see Figure 12, and e.g. 1.7 of [39]). Pulling back the orientation of M with the requirement that the normal disks have $+1$ intersection with W means the Whitney section determines a well-defined (independent of the orientation of W) relative Euler number $\omega(W) \in \mathbb{Z}$ which represents the obstruction to extending the Whitney section across W . Following traditional terminology, when $\omega(W)$ vanishes W is said to be *framed*. (Since $D^2 \times D^2$ has a unique trivialization up to homotopy, this terminology is only mildly abusive.) In general when $\omega(W) = k$, we say that W is *k -twisted*, or just *twisted* if the value of $\omega(W)$ is not specified. So a 0-twisted Whitney disk is a framed Whitney disk.




Figure 12: The Whitney section over the boundary of a framed Whitney disk is indicated by the dotted loop shown on the left for an embedded Whitney disk W in a 3-dimensional slice of 4-space. On the right is shown an embedding into 3-space of the normal disk-bundle over ∂W , indicating how the Whitney section determines a well-defined nowhere vanishing section which lies in the I -sheet and is normal to the J -sheet.

Note that a *framing* of ∂A_i (respectively A_i) is by definition a trivialization of the normal bundle of the immersion. If the ambient 4-manifold is oriented, this is equivalent to an orientation and a nonvanishing normal vector field on ∂A_i (respectively A_i). The twisting $\omega(A_i) \in \mathbb{Z}$ of an order zero surface is also defined when a framing of ∂A_i is given, and A_i is said to be *framed* when $\omega(A_i) = 0$.

2.3 Twisted Whitney towers and their orientations

In the definition of an order n Whitney tower given just above (following [9, 34, 35, 38]) all Whitney disks and order zero surfaces are required to be framed. It turns out that the natural generalization to twisted Whitney towers involves allowing twisted Whitney disks only in at least “half the order” as follows:

Definition 2.4 A *twisted Whitney tower of order 0* is a collection of properly immersed surfaces in a 4-manifold (without any framing requirement).

For $n > 0$, a *twisted Whitney tower of order $(2n - 1)$* is just a (framed) Whitney tower of order $(2n - 1)$ as in Definition 2.3 above.

For $n > 0$, a *twisted Whitney tower of order $2n$* is a Whitney tower having all intersections of order less than $2n$ paired by Whitney disks, with all Whitney disks of order less than n required to be framed, but Whitney disks of order at least n allowed to be twisted.

Remark 2.2 Note that, for any n , an order n (framed) Whitney tower is also an order n twisted Whitney tower. We may sometimes refer to a Whitney tower as a *framed* Whitney tower to emphasize the distinction, and will always use the adjective “twisted” in the setting of Definition 2.4.

Remark 2.3 The convention of allowing only order $\geq n$ twisted Whitney disks in order $2n$ twisted Whitney towers is explained both algebraically and geometrically in [11] via relationships with Milnor invariants and grope duality. In any event, an order $2n$ twisted Whitney tower can always be modified so that all its Whitney disks of order $> n$ are framed (see **notation and conventions** in section 4.1), so the twisted Whitney disks of order equal to n are the important ones (i.e. they may represent obstructions to “raising the order”).

Whitney tower orientations

Orientations on order zero surfaces in a Whitney tower \mathcal{W} are fixed, and required to induce the orientations on their boundaries. After choosing and fixing orientations on all the Whitney disks in \mathcal{W} , the associated trees are embedded in \mathcal{W} so that the vertex orientations are induced from the Whitney disk orientations, with the descending edges of each trivalent vertex enclosing the *negative intersection point* of the corresponding Whitney disk, as in Figure 11. (In fact, if a tree t has more than one trivalent vertex which corresponds to the same Whitney disk, then t will only be immersed in \mathcal{W} , but this immersion can be taken to be a local embedding around each trivalent vertex of t as in Figure 11.)

This “negative corner” convention, which differs from the positive corner convention in the older papers [9, 38], is compatible with standard orientation and commutator conventions for group elements used in [11]. With these conventions, different choices of orientations on Whitney disks in \mathcal{W} correspond to antisymmetry relations (as explained in [38]).

2.4 Intersection invariants for (twisted) Whitney towers

The obstruction theory of [38] in the current simply connected setting works as follows.

Definition 2.5 The *order n intersection invariant* $\tau_n(\mathcal{W})$ of an order n Whitney tower \mathcal{W} is defined to be

$$\tau_n(\mathcal{W}) := \sum \epsilon_p \cdot t_p \in \mathcal{T}_n$$

where the sum is over all order n intersections p , with $\epsilon_p = \pm 1$ the usual sign of a transverse intersection point.

As mentioned in the introduction, if L bounds $\mathcal{W} \subset B^4$ with $\tau_n(\mathcal{W}) = 0 \in \mathcal{T}_n$, then L bounds a Whitney tower of order $n + 1$. This is a special case of the simply connected version of the more general Theorem 2 of [38]. We will use the following version of Theorem 2 of [38] where the order zero surfaces are either properly immersed disks in B^4 or properly immersed annuli in $S^3 \times I$:

Theorem 2.6 ([38]) *If a collection A of properly immersed surfaces in a simply connected 4-manifold supports an order n Whitney tower \mathcal{W} with $\tau_n(\mathcal{W}) = 0 \in \mathcal{T}_n$, then A is regularly homotopic (rel ∂) to A' which supports an order $n + 1$ Whitney tower.*

In Section 4.4 this theorem will be strengthened by showing that the same conclusion holds if $\tau_n(\mathcal{W})$ vanishes in the reduced group $\tilde{\mathcal{T}}_n$ (Definition 1.6).

The intersection invariants for Whitney towers are extended to twisted Whitney towers as follows:

Definition 2.7 The abelian group $\mathcal{T}_{2n-1}^\infty$ is the quotient of \mathcal{T}_{2n-1} by the *boundary-twist relations*:

$$\langle (i, J), J \rangle = i \prec_J = 0$$

Here J ranges over all order $n - 1$ rooted trees.

The boundary-twist relations correspond geometrically to the fact that performing a boundary twist (Figure 18) on an order n Whitney disk $W_{(i,J)}$ creates an order $2n - 1$ intersection point $p \in W_{(i,J)} \cap W_J$ with associated tree $t_p = \langle (i, J), J \rangle$ (which is 2-torsion by the AS relations) and changes $\omega(W_{(i,J)})$ by ± 1 . Since order n twisted Whitney disks are allowed in an order $2n$ Whitney tower such trees do not represent obstructions to the existence of the next order twisted tower.

For any rooted tree J we define the corresponding ∞ -tree, denoted by J^∞ , by labeling the root univalent vertex with the symbol “ ∞ ”:

$$J^\infty := \infty - J$$

Definition 2.8 The abelian group \mathcal{T}_{2n}^∞ is the free abelian group on order $2n$ trees and order n ∞ -trees, modulo the following relations:

- (i) AS and IHX relations on order $2n$ trees (Figure 5)

- (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: $2 \cdot J^\infty = \langle J, J \rangle$

Here the AS and IHX relations are as usual, but they only apply to non- ∞ trees. The *symmetry relation* corresponds to the fact that the relative Euler number $\omega(W)$ 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 below in Lemma 4.1. The *interior-twist relation* corresponds to the fact that creating a ± 1 self-intersection in a W_J changes the twisting by ∓ 2 (Figure 19).

Remark 2.4 The symmetry, twisted IHX, and interior twist relations in \mathcal{T}_{2n}^∞ have a surprisingly natural algebraic interpretation that we explain in [15]. The idea is to extend the map $J \mapsto J^\infty$ to a *symmetric quadratic refinement* q of the bilinear form $\langle \cdot, \cdot \rangle$ on the free quasi-Lie algebra of rooted trees (the intersection form on Whitney disks) by defining $q(J) = J^\infty$ and extending to linear combinations by the formula

$$q(J + K) := J^\infty + K^\infty + \langle J, K \rangle$$

Expanding $q(I - H + X) = 0$ leads to the 6-term IHX relation

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

which is equivalent to the twisted IHX relation in the presence of the interior-twist relations. Those in turn follow by setting $K := -J$ from the symmetry relation. In [15] we show that \mathcal{T} is the universal home for invariant symmetric bilinear forms on free quasi-Lie algebras, and that \mathcal{T}_{2k}^∞ is the universal (symmetric quadratic) refinement of this form in order k .

Remark 2.5 We discovered in [9] that the (framed) IHX relation can be realized in three dimensions as well as four, and it is interesting to note that many of the relations that we obtain for twisted Whitney towers in four dimensions can also be realized by rooted clasper surgeries (grope cobordisms) in three dimensions. Here the twisted Whitney disk corresponds to a ± 1 framed leaf of a clasper. For example the relation $I^\infty = H^\infty + X^\infty - \langle H, X \rangle$ has the following clasper explanation. I^∞ represents a clasper with one isolated twisted leaf. By the topological IHX relation, one can replace I^∞ by two clasplers of the form H^∞ and $(-X)^\infty = X^\infty$ embedded in a regular neighborhood of

the original clasper with leaves parallel to the leaves of the original. The twisted leaves are now linked together, so applying Habiro's zip construction (which complicates the picture considerably) one gets three tree claspers, of the form H^∞ , X^∞ and $\langle H, -X \rangle$ respectively.

Similarly, the relation $2 \cdot J^\infty = \langle J, J \rangle$ has an interpretation where one takes a clasper which represents J^∞ and splits off a geometrically canceling parallel copy, representing the tree J^∞ . Again, because the twisted leaves link, we also get the term $\langle J, -J \rangle$.

These observations will be enlarged upon in [14] to analyze filtrations on homology cylinders and string links.

Recall from Definition 2.4 (and Remark 2.3) that twisted Whitney disks only occur in even order twisted Whitney towers, and only those of half-order are relevant to the obstruction theory.

Definition 2.9 The *order n intersection intersection invariant* $\tau_n^\infty(\mathcal{W})$ of an order n twisted Whitney tower \mathcal{W} is defined to be

$$\tau_n^\infty(\mathcal{W}) := \sum \epsilon_p \cdot t_p + \sum \omega(W_J) \cdot J^\infty \in \mathcal{T}_n^\infty$$

where the first sum is over all order n intersections p and the second sum is over all order $n/2$ Whitney disks W_J with twisting $\omega(W_J) \in \mathbb{Z}$. For $n = 0$, recall from 2.1 above our notational convention that W_j denotes A_j , and that $\omega(A_j) \in \mathbb{Z}$ is the relative Euler number of the normal bundle of A_j with respect to the given framing of ∂A_j as in 2.2 .

By splitting the twisted Whitney disks, as explained in subsection 2.5 below, for $n > 0$ we may actually assume that all non-zero $\omega(W_J) \in \{\pm 1\}$, just like the signs ϵ_p .

As in the framed case, the vanishing of τ_n^∞ is sufficient for the existence of a twisted Whitney tower of order $(n + 1)$, and the proof in Section 4 of Theorem 1.8 (describing the twisted realization maps $R_n^\infty : \mathcal{T}_n^\infty \rightarrow \mathcal{W}_n^\infty$) will be based on the following analogue of the framed order-raising Theorem 2.6 to the twisted setting:

Theorem 2.10 *If a collection A of properly immersed surfaces in a simply connected 4-manifold supports an order n twisted Whitney tower \mathcal{W} with $\tau_n^\infty(\mathcal{W}) = 0 \in \mathcal{T}_n^\infty$, then A is regularly homotopic (rel ∂) to A' which supports an order $n + 1$ twisted Whitney tower.*

The proof of Theorem 2.10 is given in Section 4 below.

Proofs of the “order-raising” Theorems 2.10 and 2.6 (and its strengthening Theorem 4.4 below) depend on realizing the relations in the target groups by controlled manipulations of Whitney towers. The next two subsections introduce combinatorial notions useful for describing the algebraic effect of such geometric constructions.

For the rest of this section we assume our Whitney towers are of positive order for convenience of notation.

Intersection forests

Recall that the trees associated to intersections and Whitney disks in a Whitney tower can be considered to be immersed in the Whitney tower, with vertex orientations induced by the Whitney tower orientation, as in Figure 11.

Definition 2.11 The *intersection forest* $t(\mathcal{W})$ of a framed Whitney tower \mathcal{W} is the disjoint union of signed trees associated to all unpaired intersections p in \mathcal{W} :

$$t(\mathcal{W}) = \coprod_p \epsilon_p \cdot t_p$$

with ϵ_p the sign of the intersection point p . For \mathcal{W} of order n , we can think of the signed order n trees in $t(\mathcal{W})$ as an “abelian word” in the generators $\pm t_p$ which represents $\tau_n(\mathcal{W}) \in \mathcal{T}_n$. More precisely, $t(\mathcal{W})$ is an element of the free abelian monoid, with unit \emptyset , generated by (isomorphism classes of) signed trees, trivalent, labeled and vertex-oriented as usual. We emphasize that there are no cancellations or other relations here.

Remark 2.6 In the older papers [9, 34, 38] we referred to $t(\mathcal{W})$ as the “geometric intersection tree” (and to the group element $\tau_n(\mathcal{W})$ as the order n intersection “tree”, rather than “invariant”), but the term “forest” better describes the disjoint union of (signed) trees $t(\mathcal{W})$.

Similarly to the framed case, the *intersection forest* $t(\mathcal{W})$ of a *twisted* Whitney tower \mathcal{W} is the disjoint union of signed trees associated to all unpaired intersections p in \mathcal{W} and integer-coefficient ∞ -trees associated to all non-trivially twisted Whitney disks W_J in \mathcal{W} :

$$t(\mathcal{W}) = \coprod_p \epsilon_p \cdot t_p + \coprod_J \omega(W_J) \cdot J^\infty$$

with $\omega(W_J) \in \mathbb{Z}$ the twisting of W_J . Again, there are no cancellations or relations (and the informal “+” sign in the expression is purely cosmetic).

We will see in the next subsection that all the trees can be made to be disjoint in \mathcal{W} , with all non-zero $\omega(W_J) = \pm 1$, so that $t(\mathcal{W})$ is also a topological disjoint union which corresponds to an element in the free abelian monoid generated by (isomorphism classes of) signed trees, and signed ∞ -trees.

2.5 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 [38] (Lemma 3.5 of [34]), by performing finger moves along Whitney disks guided by arcs connecting the Whitney disk boundary arcs (see Figure 3). Implicit in this construction is that the finger moves preserve the Whitney disk framings (by not twisting relative to the Whitney disk that is being split – see Figure 13). 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 will simplify the combinatorics to use “twisted” finger moves to similarly split-off 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).

Lemma 2.12 *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 disjoint union of non- ∞ trees $\coprod_p \epsilon_p \cdot t_p \subset t(\mathcal{W})$ is isomorphic to the disjoint union of non- ∞ trees $\coprod_{p'} \epsilon_{p'} \cdot t_{p'} \subset t(\mathcal{W}')$.*
- (ii) *Each $\omega(W_J) \cdot J^\infty$ in $t(\mathcal{W})$ gives rise to the disjoint union of exactly $|\omega(W_J)|$ -many $\pm 1 \cdot J^\infty$ in \mathcal{W}' , where the sign \pm corresponds to the sign of $\omega(W_J)$.*

Proof Illustrated in Figure 13 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 13 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




Figure 13: 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.)

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 of the framed Whitney disks as in Lemma 13 of [38] gives the result. \square

3 The realization maps

This section contains clarifications and proofs of Theorems 1.3 and 1.8 from the introduction which state the existence of surjections $R_n: \mathcal{T}_n \rightarrow \mathcal{W}_n$ and $R_n^\infty: \mathcal{T}_n^\infty \rightarrow \mathcal{W}_n^\infty$ for all n , in particular exhibiting the sets \mathcal{W}_n and \mathcal{W}_n^∞ as finitely generated abelian groups under connected sum.

All proofs in this section apply in the reduced setting as well, and the constructions described here also define the surjections $\tilde{R}_n: \tilde{\mathcal{T}}_n \rightarrow \mathcal{W}_n$ described in the introduction.

Recall that our manifolds are assumed oriented, but orientations are suppressed from the discussion as much as possible. In the following an orientation is fixed once and for all on S^3 ; and a *framed link* has oriented components, each equipped with a nowhere-vanishing normal section.

Definition 3.1 A framed link $L \subset S^3 = \partial B^4$ bounds an order n Whitney tower \mathcal{W} if $\mathcal{W} \subset B^4$ is an order n Whitney tower whose order zero surfaces are immersed disks bounded by the components of L , as in Definition 2.3.

Similarly, a framed 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 , as in Definition 2.4.

Definition 3.2 For $n \geq 1$, framed links L_0 and L_1 in S^3 are *Whitney tower concordant of order n* if the i th components of $L_0 \subset S^3 \times \{0\}$ and $-L_1 \subset S^3 \times \{1\}$ cobound an immersed annulus A_i for each i such that the A_i are transverse and support an order n Whitney tower. If the A_i support a *twisted* order n Whitney tower then L_0 and L_1 are said to be *twisted Whitney tower concordant of order n* .

Note that a (twisted) Whitney tower concordance preserves framings on L_0 and L_1 (as links in S^3) since, for all i , $\omega(A_i) = 0$ because all self-intersections of the A_i come in (geometrically) canceling pairs in any (twisted) Whitney tower of order $n \geq 1$.

Recall from the introduction that the set of m -component framed links in S^3 which bound order n (twisted) Whitney towers in B^4 is denoted by $\mathbb{W}_n = \mathbb{W}_n(m)$ (resp. \mathbb{W}_n°); and the quotient of \mathbb{W}_n by the equivalence relation of order $n + 1$ (twisted) Whitney tower concordance is denoted by W_n (resp. W_n°).

Throughout this section the twisted setting mirrors the framed setting, with discussions and arguments given simultaneously.

We begin by deriving from the “order-raising” Theorem 2.6 the following essential criterion for links to represent equal elements in the associated graded W_n :

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

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

Conversely, suppose that L_0 and L_1 bound order n Whitney towers \mathcal{W}_0 and \mathcal{W}_1 in 4–balls B_0^4 and B_1^4 , with $\tau_n(\mathcal{W}_0) = \tau_n(\mathcal{W}_1)$. 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 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 $\tau_n(\mathcal{W}_1)$, it follows that \mathcal{V} has vanishing order n intersection invariant:

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

So by Theorem 2.6, A is homotopic (rel ∂) to A' supporting an order $n + 1$ Whitney tower, and hence L_0 and L_1 are equal in \mathcal{W}_n . \square

Remark 3.1 The analogous statement and proof of Corollary 3.3 holds in the twisted case (with the “twisted order-raising” Theorem 2.10 playing the role of Theorem 2.6). For this case, we’ll spell out the statement carefully but in several instances below we will just state that the twisted case is analogous: Links L_0 and L_1 in \mathbb{W}_n^∞ represent the same element of \mathcal{W}_n^∞ if and only if there exist order n twisted Whitney towers \mathcal{W}_0 and \mathcal{W}_1 in B^4 bounded by L_0 and L_1 respectively such that $\tau_n^\infty(\mathcal{W}_0) = \tau_n^\infty(\mathcal{W}_1) \in \mathcal{T}_n^\infty$.

Remark 3.2 Remark 3.1 similarly applies to the reduced setting by Theorem 4.4 below, although we will omit further reference to $\tilde{\mathcal{T}}$ in this section.

3.1 Band sums of links

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 \mathcal{W}_n (or in \mathcal{W}_n^∞) does not depend on β :

Lemma 3.4 *For links L and L' representing elements of \mathcal{W}_n , any band sum $L\#_\beta L'$ represents an element of \mathcal{W}_n which only depends on the equivalence classes of L and L' in \mathcal{W}_n . The same statement holds in \mathcal{W}_n^∞ .*

Proof We shall only give the proof in the framed case, the twisted case is analogous. If L_0 and L_1 represent the same element of \mathcal{W}_n , and if L'_0 and L'_1 represent the same element of \mathcal{W}_n , then by Corollary 3.3 above, for $i = 0, 1$, there are order n Whitney towers \mathcal{W}_i and \mathcal{W}'_i bounding L_i and L'_i such that $\tau_n(\mathcal{W}_0) = \tau_n(\mathcal{W}_1)$ and $\tau_n(\mathcal{W}'_0) = \tau_n(\mathcal{W}'_1)$. By Lemma 3.5 just below, $L_i\#_{\beta_i} L'_i$ bounds $\mathcal{W}_i^\#$ for $i = 0, 1$, with

$$\tau_n(\mathcal{W}_0^\#) = \tau_n(\mathcal{W}_0) + \tau_n(\mathcal{W}'_0) = \tau_n(\mathcal{W}_1) + \tau_n(\mathcal{W}'_1) = \tau_n(\mathcal{W}_1^\#)$$

so again by Corollary 3.3, $L_0 \#_{\beta_0} L'_0$ is order $n+1$ Whitney tower concordant to $L_1 \#_{\beta_1} L'_1$, hence $L_0 \#_{\beta_0} L'_0$ and $L_1 \#_{\beta_1} L'_1$ represent the same element of W_n . \square

Lemma 3.5 *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 $t(\mathcal{W}^\#) = t(\mathcal{W}) \amalg t(\mathcal{W}')$, where $t(\mathcal{V})$ denotes the intersection forest of a Whitney tower \mathcal{V} as above in subsection 2.4.*

Proof Let B and B' be the 3-balls in the link complements used to form the S^3 containing $L \#_{\beta} L'$. Then gluing together the two 4-balls containing \mathcal{W} and \mathcal{W}' along B and B' forms B^4 containing $L \#_{\beta} L'$ in its boundary. Take $\mathcal{W}^\#$ to be the boundary band sum of \mathcal{W} and \mathcal{W}' along the order zero disks guided by the bands β , with the interiors of the bands perturbed slightly into the interior of B^4 . It is clear that $t(\mathcal{W}^\#)$ is just the disjoint union $t(\mathcal{W}) \amalg t(\mathcal{W}')$ since no new singularities have been created. \square

3.2 Definition of the realization maps

The realization maps R_n are defined as follows: Given any group element $g \in \mathcal{T}_n$, by Lemma 3.6 just below there exists an m -component link $L \subset S^3$ bounding an order n Whitney tower $\mathcal{W} \subset B^4$ such that $\tau_n(\mathcal{W}) = g \in \mathcal{T}_n$. Define $R_n(g)$ 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 3.3. The twisted realization map R_n^∞ is defined via Lemma 3.6 the same way using twisted Whitney towers.

Lemma 3.6 *For any disjoint union $\amalg_p \epsilon_p \cdot t_p + \amalg_J \omega(W_J) \cdot J^\infty$ there exists an m -component link L bounding a twisted Whitney tower \mathcal{W} with intersection forest $t(\mathcal{W}) = \amalg_p \epsilon_p \cdot t_p + \amalg_J \omega(W_J) \cdot J^\infty$. If the disjoint union contains no ∞ -trees then all Whitney disks in \mathcal{W} are framed.*

Note that if in the disjoint union all non- ∞ trees are order at least n and all ∞ -trees are order at least $n/2$ then \mathcal{W} will have order n .

Proof It suffices to consider the cases where the disjoint union consists of just a single (signed) tree or ∞ -tree since by Lemma 3.5 any sum of such trees can then be realized by band sums of links.

The following algorithm, in the untwisted case, is the algorithm called "Bing-doubling along a tree" by Cochran and used in Section 7 of [2] and Theorem 3.3 of [3] to produce links in S^3 with prescribed (first non-vanishing) Milnor invariants.




Figure 14: Pushing into B^4 from left to right: A Hopf link in $S^3 = \partial B^4$ bounds embedded disks $D_1 \cup D_2 \subset B^4$ which intersect in a point p , with $t_p = \langle 1, 2 \rangle$.




Figure 15: Pushing into B^4 from left to right: The disks D_{i_2} and D_j extend to the right-most picture where they are completed by capping off the unlink. The disk D_{i_1} only extends to the middle picture where the intersections between D_{i_1} and D_{i_2} are paired by the Whitney disk $W_{(i_1, i_2)}$, that has a single interior intersection $p \in W_{(i_1, i_2)} \cap D_j$ with $t_p = \langle (i_1, i_2), j \rangle$.

Realizing order zero trees and ∞ -trees. A 0-framed Hopf link bounds an order zero Whitney tower $\mathcal{W} = D_1 \cup D_2 \subset B^4$, where the two embedded disks D_1 and D_2 have a single interior intersection point p with $t_p = \langle 1, 2 \rangle = 1 \text{ --- } 2$ (see Figure 14). Assuming appropriate fixed orientations of B^4 and S^3 , the sign ϵ_p associated to p is the usual sign of the Hopf link. So taking a 0-framed $(m - 2)$ -component trivial link together with a Hopf link (as the i th and j th components) gives an m -component link L bounding \mathcal{W} with $t(\mathcal{W}) = \epsilon_p \cdot \langle i, j \rangle = \epsilon_p \cdot i \text{ --- } j$, for any $\epsilon_p = \pm 1$, and $i \neq j$.

To realize the tree $\pm i \text{ --- } i$, we can use the unlink with framings 0, except that the component labeled by the index i has framing ± 2 . Similarly, if the component has framing ± 1 then the resulting tree is $\pm \infty \text{ --- } i$.

Realizing order 1 trees. Consider now a link L whose i th and j th components form a Hopf link $L^i \cup L^j$ bounding disks $D_i \cup D_j \subset B^4$ with transverse intersection $p = D_i \cap D_j$. Assume that $D_i \cup D_j$ extends to an order zero Whitney tower \mathcal{W} bounded by L with $t(\mathcal{W}) = \epsilon_p \cdot t_p = \epsilon_p \cdot \langle i, j \rangle$.




Figure 16: Pushing into B^4 from left to right: An i - and j -labeled n -twisted Bing-double (case $n = 2$) of the unknot in $S^3 = \partial B^4$ bounds disks D_i and D_j whose intersections are paired by a Whitney disk $W_{(i,j)}$. D_j extends to the right-hand picture but D_i only extends to the middle picture, where the boundary of $W_{(i,j)}$ is indicated by the dark arcs. The rest of $W_{(i,j)}$ extends into the right-hand picture where disjointly embedded disks bounded by the unlink complete both $W_{(i,j)}$ and D_j . The interior of $W_{(i,j)}$ is embedded and disjoint from both D_i and D_j . Figure 17 shows that $W_{(i,j)}$ is twisted, with $\omega(W_{(i,j)}) = n$.

Replacing L^i by an untwisted Bing-double $L^{i_1} \cup L^{i_2}$ results in a new sublink of Borromean rings $L^{i_1} \cup L^{i_2} \cup L^j$ bounding disks $D_{i_1} \cup D_{i_2} \cup D_j$ as indicated in Figure 15, with D_{i_1} and D_{i_2} intersecting in a canceling pair of intersections paired by an order 1 Whitney disk $W_{(i_1, i_2)}$, which can be formed from D_{i_1} with a small collar removed, so that $W_{(i_1, i_2)}$ has a single intersection with D_j corresponding to the original $p = D_i \cap D_j$. (One can think of D_{i_1} and D_{i_2} as being formed by the trace of the obvious pulling-apart homotopy that shrinks L^{i_1} and L^{i_2} down in a tubular neighborhood of L^i , with the canceling pair of intersections between D_{i_1} and D_{i_2} being created as the clasps are pulled apart.)

The effect of this Bing-doubling operation on the intersection forest is that the original order zero $t_p = \langle i, j \rangle$ has given rise to the order 1 tree $\langle \langle i_1, i_2 \rangle, j \rangle$. Switching the orientation on one of the new components changes the sign of p , as can be checked using our orientation conventions. By relabeling and/or banding together components of this new link any labels on this order 1 tree can be realized. Since the doubling was untwisted, $W_{(i_1, i_2)}$ is framed (see Figures 16 and 17), so the Whitney tower bounded by the new link is order 1.

Realizing order n trees. Since any order n tree can be gotten from some order $n - 1$ tree by attaching two new edges to a univalent vertex as in the previous paragraph, it follows inductively that any order n tree is the intersection forest of a Whitney tower bounded by some link. (First create a distinctly-labeled tree of the desired ‘shape’ by doubling, then correct the labels by interior band-summing.)

Realizing ∞ -trees of order 1. As illustrated (for the case $n = 2$) in Figures 16 and 17, the n -twisted Bing-double of the unknot (with components labeled i and j) bounds




Figure 17: The Whitney section over $\partial W_{(i,j)}$ (from Figure 16) is indicated by the dashed arcs on the left. The twisting $\omega(W_{(i,j)}) = n$ (the obstruction to extending the Whitney section across the Whitney disk) corresponds to the n -twisting of the Bing-doubling operation.

an order 2 twisted Whitney tower \mathcal{W} with $t(\mathcal{W}) = n \cdot (i, j)^\infty = n \cdot \infty \prec_j^i$. Banding together the two components would yield a knot realizing $(i, i)^\infty$.

Realizing ∞ -trees of order n . By applying iterated untwisted Bing-doubling operations to the i - and j -labeled components of the order 1 case, one can construct for any rooted tree (I, J) a link bounding a twisted Whitney tower \mathcal{W} with $t(\mathcal{W}) = n \cdot (I, J)^\infty$. For instance, if in the construction of Figure 16 the j -labeled link component is replaced by an untwisted Bing-double, then the disk D_j in that construction would be replaced by a (framed) Whitney disk $W_{(j_1, j_2)}$, and the n -twisted $W_{(i, j)}$ would be replaced by an n -twisted $W_{(i, (j_1, j_2))}$. (As for non- ∞ trees above, first create a distinctly-labeled tree of the desired ‘shape’ by doubling, then correct the labels by interior band-summing.) \square

3.3 Surjectivity of the realization maps

In this section we will prove Theorems 1.3 and 1.8: The realization maps R_n and R_n^∞ are epimorphisms.

We will show, moreover, that \mathbb{W}_n is the set of framed links $L \in \mathbb{W}_n$ modulo the relation that $[L_1] = [L_2] \in \mathbb{W}_n$ if and only if $L_1 \# -L_2$ lies in \mathbb{W}_{n+1} , where $-L$ is the mirror image of L with reversed framing.

Proof From Lemma 3.4 the band sum of links gives well-defined operations in \mathbb{W}_n and \mathbb{W}_n^∞ which are clearly associative and commutative, with the m -component unlink representing an identity element. The realization maps are homomorphisms by Lemma 3.5 and surjectivity is proven as follows: Given any link $L \in \mathbb{W}_n$, choose a Whitney tower \mathcal{W} of order n with boundary L and compute $\tau := \tau_n(\mathcal{W})$. Then take $L' := R_n(\tau)$, a link that’s obviously in the image of R_n and for which we know a

Whitney tower \mathcal{W}' with boundary L' and $\tau(\mathcal{W}') = \tau$. By Corollary 3.3 it follows that L and L' represent the same element in W_n .

If L_0 and L_1 represent the same element of W_n (resp. W_n^∞), then by Corollary 3.3 there exist order n (twisted) Whitney towers \mathcal{W}_0 and \mathcal{W}_1 in B^4 bounded by L_0 and L_1 respectively such that $\tau_n(\mathcal{W}_0) = \tau_n(\mathcal{W}_1) \in \mathcal{T}_n$ (resp. $\tau_n^\infty(\mathcal{W}_0) = \tau_n^\infty(\mathcal{W}_1) \in \mathcal{T}_n^\infty$). We want to show that $L_0 \# -L_1$ bounds an order $n + 1$ (twisted) Whitney tower, which will follow from Lemma 3.5 and the ‘‘order-raising’’ Theorem 2.6 (respectively Theorem 2.10) if $-L_1$ bounds an order n (twisted) Whitney tower $\overline{\mathcal{W}}_1$ such that $\tau_n(\overline{\mathcal{W}}_1) = -\tau_n(\mathcal{W}_1) \in \mathcal{T}_n$ (resp. $\tau_n^\infty(\overline{\mathcal{W}}_1) = -\tau_n^\infty(\mathcal{W}_1) \in \mathcal{T}_n^\infty$). If r denotes the reflection on S^3 which sends L_1 to $-L_1$, then the product $r \times \text{id}$ of r with the identity is an involution on $S^3 \times I$, and the image $r \times \text{id}(\mathcal{W}_1)$ of \mathcal{W}_1 is such a $\overline{\mathcal{W}}_1$. To see this, observe that $r \times \text{id}$ switches the signs of all transverse intersection points, and is an isomorphism on the oriented trees in \mathcal{W}_1 ; and hence switches the signs of all Whitney disk framing obstructions (which can be computed as intersection numbers between Whitney disks and their push-offs) – note that $r \times \text{id}$ is only being applied to \mathcal{W}_1 , while $S^3 \times I$ is fixed.

Assume now that $L_0 \# -L_1 \subset S^3$ bounds an order $n + 1$ (twisted) Whitney tower $\mathcal{W} \subset B^4$. By the definition of connected sum, S^3 decomposes as the union of two disjoint 3–balls B_0 and B_1 containing L_0 and $-L_1$, joined together by the $S^2 \times I$ through which passes the bands guiding the connected sum. Taking another 4–ball with the same decomposition of its boundary 3–sphere, and gluing the 4–balls together by identifying the boundary 2–spheres of the 3–balls, and identifying the $S^2 \times I$ subsets by the identity map, forms $S^3 \times I$ containing an order $n + 1$ (twisted) Whitney tower concordance between L_0 and $-L_1$ which consists of \mathcal{W} together with the parts of the connected-sum bands that are contained in $S^2 \times I$. \square

4 Implications of the twisted IHX construction

This section is mostly dedicated to proving the ‘‘twisted order-raising’’ Theorem 2.10 of Section 2, which was used in Section 3 to construct the twisted realization maps. A key step in the proof given in section 4.1 involves a geometric realization of the twisted IHX relation as described in Lemma 4.1 below.

In section 4.3, Lemma 4.1 is also used to show how any order $2n$ twisted Whitney tower can be converted into an order $2n - 1$ framed Whitney tower. This result (Lemma 4.2) will be used later in the proof of Theorem 5.1 in Section 5.4.

Then in section 4.4, Lemma 4.1 is used again to prove the order-raising Theorem 4.4 in the reduced setting.

Recall the statement of Theorem 2.10: If a collection A of properly immersed surfaces in a simply connected 4–manifold supports an order n twisted Whitney tower \mathcal{W} with $\tau_n^\infty(\mathcal{W}) = 0 \in \mathcal{T}_n^\infty$, then A is regularly homotopic (rel ∂) to A' supporting an order $n + 1$ twisted Whitney tower.

Recall also from Definition 2.11 that the intersection forest $t(\mathcal{W})$ of an order n twisted Whitney tower \mathcal{W} is a disjoint union of signed trees which can be considered to be immersed in \mathcal{W} . The order n trees in $t(\mathcal{W})$ (together with the order $n/2$ ∞ -trees if n is even) represent $\tau_n^\infty(\mathcal{W}) \in \mathcal{T}_n^\infty$, and the proof of Theorem 2.10 involves controlled manipulations of \mathcal{W} which first convert $t(\mathcal{W})$ into “algebraically canceling” pairs of isomorphic trees with opposite signs, and then exchange these for “geometrically canceling” intersection points which are paired by a new layer of Whitney disks. We pause here to clarify these notions:

Algebraic versus geometric cancellation: Note that $t(\mathcal{W})$ is a combinatorial object which by Lemma 2.12 above can be considered geometrically as the image of an embedding in \mathcal{W} . If Whitney disks W_I and W_J in \mathcal{W} intersect transversely in a pair of points p and p' , then t_p and $t_{p'}$ are isomorphic (as labeled, oriented trees). If p and p' have opposite signs, and if the ambient 4–manifold is simply connected, then there exists a Whitney disk $W_{(I,J)}$ pairing p and p' , and we say that $\{p, p'\}$ is a *geometrically canceling pair*. In this setting we also refer to $\{\epsilon_p \cdot t_p, \epsilon_{p'} \cdot t_{p'}\}$ as a geometrically canceling pair of signed trees in $t(\mathcal{W})$ (regarding them as subsets of \mathcal{W} associated to the geometrically canceling pair of points).

On the other hand, given transverse intersections p and p' in \mathcal{W} with $t_p = t_{p'}$ (as labeled oriented trees) and $\epsilon_p = -\epsilon_{p'}$, we say that $\{p, p'\}$ is an *algebraically canceling pair* of intersections, and similarly call $\{\epsilon_p \cdot t_p, \epsilon_{p'} \cdot t_{p'}\}$ an algebraically canceling pair of signed trees in $t(\mathcal{W})$. Changing the orientations at a *pair* of trivalent vertices in any tree t_p does not change its value in \mathcal{T} by the AS relations, and (as discussed in 3.4 of [38]) such orientation changes can be realized by changing orientations of Whitney disks in \mathcal{W} together with our orientation conventions (2.3).

Any geometrically canceling pair is also an algebraically canceling pair, but the converse is clearly not true as an algebraically canceling pair can have *corresponding trivalent vertices* lying in *different Whitney disks*. A process for converting algebraically canceling pairs into geometrically canceling pairs by manipulations of the Whitney tower is described in 4.5 and 4.8 of [38].

Similarly, if a pair of twisted Whitney disks W_{J_1} and W_{J_2} have isomorphic (unoriented) trees J_1^∞ and J_2^∞ with opposite twistings $\omega(W_{J_1}) = -\omega(W_{J_2})$, then the Whitney disks form an *algebraically canceling* pair (as do the corresponding signed ∞ -trees in $t(\mathcal{W})$). Note that the orientations of the ∞ -trees are not relevant here by the independence of $\omega(W)$ from the orientation of W and the symmetry relations in \mathcal{T}^∞ . A geometric construction for eliminating algebraically canceling pairs of twisted Whitney disks from a twisted Whitney tower will be described below.

4.1 Proof of the twisted order-raising Theorem 2.10

To motivate the proof we summarize here how the methods of [9, 34, 38] (as described in Section 4 of [38]) apply in the framed setting to prove the analogous order-raising theorem in framed setting (Theorem 2.6 of Section 2): The first part of the proof changes the intersection forest $t(\mathcal{W})$ so that all trees occur in algebraically canceling pairs by using the 4-dimensional IHX construction of [9] to realize IHX relators, and by adjusting Whitney disk orientations as necessary to realize AS relations. The second part of the proof uses the Whitney move IHX construction of [34] to “simplify” the shape of the algebraically canceling pairs of trees. Then the third part of the proof uses controlled homotopies to exchange the simple algebraic canceling pairs for geometrically canceling intersection points which are paired by a new layer of Whitney disks as described in 4.5 of [38]. All constructions only change the order 0 surfaces by regular homotopies consisting of finger moves, Whitney moves, and isotopies.

Extending these methods to the present twisted setting will require two variations: realizing the new relators in \mathcal{T}_n^∞ , and achieving an analogous geometric cancellation for twisted Whitney disks corresponding to algebraically canceling pairs of (simple) ∞ -trees. We will concentrate on these new variations, referring the reader to [9, 34, 38] for the other parts just mentioned.

Notation and conventions:

By Lemma 2.12 it may be assumed that \mathcal{W} is split at each stage of the constructions throughout the proof, so that all trees in $t(\mathcal{W})$ are embedded in \mathcal{W} . In spite of modifications, \mathcal{W} will not be renamed during the proof. Throughout this section we will notate elements of $t(\mathcal{W})$ as formal sums, representing disjoint union by juxtaposition.

Note that if \mathcal{W} is an order n twisted Whitney tower, then the intersection forest $t(\mathcal{W})$ may contain higher order trees and ∞ -trees in addition to those representing $\tau_n^\infty(\mathcal{W}) \in \mathcal{T}_n^\infty$. These higher-order elements of $t(\mathcal{W})$ can be ignored throughout the proof for the following reasons: On the one hand, in a split \mathcal{W} all the constructions




Figure 18: Boundary-twisting a Whitney disk W changes $\omega(W)$ by ± 1 and creates an intersection point with one of the sheets paired by W . The horizontal arcs trace out part of the sheet, the dark non-horizontal arcs trace out the newly twisted part of a collar of W , and the grey arcs indicate part of the Whitney section over W . The bottom-most intersection in the middle picture corresponds to the ± 1 -twisting created by the move.

leading to the elimination of unpaired order n intersections (and twisted Whitney disks of order $n/2$) of \mathcal{W} can be carried out away from any higher-order elements of $t(\mathcal{W})$. Alternatively, one could first exchange all twisted Whitney disks of order greater than $n/2$ for unpaired intersections of order greater than n by boundary-twisting (Figure 18). Then, all intersections of order greater than n can be converted into many algebraically canceling pairs of order n intersections by repeatedly “pushing down” unpaired intersections until they reach the order zero disks, as illustrated for instance in Figure 12 of [35] (assuming, as we may, that \mathcal{W} contains no Whitney disks of order greater than n).

Thus, we can and will assume throughout the proof that $t(\mathcal{W})$ represents $\tau_n^\infty(\mathcal{W}) \in \mathcal{T}_n^\infty$.

The odd order case: Given \mathcal{W} of order $2n - 1$ with $\tau_{2n-1}^\infty(\mathcal{W}) = 0 \in \mathcal{T}_{2n-1}^\infty$, it will suffice to modify \mathcal{W} — while only creating unpaired intersections of order at least $2n$ and twisted Whitney disks of order at least n — so that all order $2n - 1$ trees in $t(\mathcal{W})$ come in algebraically canceling pairs of trees (since by [38] the corresponding algebraically canceling pairs of order $2n - 1$ intersection points can be exchanged for geometrically canceling intersections which are paired by Whitney disks, as mentioned just above).

Since $\tau_{2n-1}^\infty(\mathcal{W}) = 0 \in \mathcal{T}_{2n-1}^\infty$, the intersection forest $t(\mathcal{W})$ is in the span of IHX and boundary-twist relators, after choosing Whitney disk orientations to realize AS relations as necessary. By locally creating intersection trees of the form $+I - H + X$ using the 4-dimensional geometric IHX theorem of [9] (and by choosing Whitney disk orientations to realize AS relations as needed), \mathcal{W} can be modified so that all order $2n - 1$ trees in $t(\mathcal{W})$ either come in algebraically canceling pairs, or are boundary-relator trees of the form $\pm\langle(i, J), J\rangle$.

For each tree of the form $t_p = \pm\langle(i, J), J\rangle$ we can create an algebraically canceling $t_{p'} = \mp\langle(i, J), J\rangle$ at the cost of only creating order n ∞ -trees as follows. First use

Lemma 14 of [38] (Lemma 3.6 of [34]) to move the unpaired intersection point p so that $p \in W_{(i,J)} \cap W_J$. Now, by boundary-twisting $W_{(i,J)}$ into its supporting Whitney disk W'_J (Figure 18), an algebraically canceling intersection $p' \in W_{(i,J)} \cap W'_J$ can be created at the cost of changing the twisting $\omega(W_{(i,J)})$ by ± 1 . Since $\langle (i, J), J \rangle$ has an order 2 symmetry, the canceling sign can always be realized by a Whitney disk orientation choice. This algebraic cancellation of t_p has been achieved at the cost of only adding to $t(\mathcal{W})$ the order n ∞ -tree $(i, J)^\infty$ corresponding to the ± 1 -twisted order n Whitney disk $W_{(i,J)}$.

Having arranged that all the order $2n - 1$ trees in $t(\mathcal{W})$ occur in algebraically canceling pairs, applying the tree-simplification and geometric cancellation described in [38] to all these algebraically canceling pairs yields an order $2n$ twisted Whitney tower \mathcal{W}' .

The even order case: For \mathcal{W} of order $2n$ with $\tau_{2n}^\infty(\mathcal{W}) = 0 \in \mathcal{T}_{2n}^\infty$, we arrange for $t(\mathcal{W})$ to consist of only algebraically canceling pairs of generators by realizing all relators in \mathcal{T}_{2n}^∞ , then construct an order $2n + 1$ twisted Whitney tower by introducing a new method for geometrically canceling the pairs of twisted Whitney disks (while the algebraically canceling pairs of non- ∞ trees lead to geometrically canceling intersections as before):

First of all, the order 0 case corresponding to linking numbers is easily checked, so we will assume $n \geq 1$. The IHX relators and AS relations for non- ∞ trees can be realized as usual. Note that any signed tree $\epsilon \cdot J^\infty \in t(\mathcal{W})$ does not depend on the orientation of the tree J because changing the orientation on the corresponding twisted Whitney disk W_J does not change $\omega(W_J)$.

For any rooted tree J the relator $\langle J, J \rangle - 2 \cdot J^\infty$ corresponding to the interior-twist relation can be realized as follows. Use finger moves to create a clean framed Whitney disk W_J . Performing a positive interior twist on W_J as in Figure 19 creates a self-intersection $p \in W_J \cap W_J$ with $t_p = \langle J, J \rangle$ and changes the twisting $\omega(W_J)$ of W_J to -2 . The negative of the relator is similarly constructed starting with a negative twist.




Figure 19: A $+1$ interior twist on a Whitney disk changes the twisting by -2 , as is seen in the pair of -1 intersections between the black vertical slice of the Whitney disk and the grey slice of a Whitney-parallel copy. Note that the pair of (positive) black-grey intersections near the $+1$ intersection is just an artifact of the immersion of the normal bundle into 4-space and does not contribute to the relative Euler number.

The relator $-I^\infty + H^\infty + X^\infty - \langle H, X \rangle$ corresponding to the twisted IHX relation is realized as follows. For any rooted tree I , create a clean framed Whitney disk W_I by finger moves. Then split this framed Whitney disk using the twisted finger move of Lemma 2.12 into two clean twisted Whitney disks with twistings $+1$ and -1 , and associated signed ∞ -trees $+I^\infty$ and $-I^\infty$, respectively. The next step is to perform a $+1$ -twisted version (described in Lemma 4.1 below) of the ‘‘Whitney move IHX’’ construction of Lemma 7.2 in [34], which will replace the $+1$ -twisted Whitney disk by two $+1$ -twisted Whitney disks having ∞ -trees $+H^\infty$ and $+X^\infty$, and containing a single negative intersection point with tree $-\langle H, X \rangle$, where H and X differ locally from I as in the usual IHX relation. Thus, any Whitney tower can be modified to create exactly the relator $-I^\infty + H^\infty + X^\infty - \langle H, X \rangle$, for any rooted tree I . The negative of the relator can be similarly realized by using Lemma 4.1 applied to the -1 -twisted I -shaped Whitney disk.

So since $\tau_{2n}^\infty(\mathcal{W})$ vanishes, it may be arranged, by realizing relators as above, that all the trees in $t(\mathcal{W})$ occur in algebraically canceling pairs. Now, by repeated applications of Lemma 4.1 below, the algebraically canceling pairs of clean ± 1 -twisted Whitney disks can be exchanged for (many) algebraically canceling pairs of clean ± 1 -twisted Whitney disks, all of whose trees are *simple* (right- or left-normed), with the ∞ -label at one end of the tree as illustrated in Figure 20 – this also creates more algebraically canceling pairs of non- ∞ trees (the ‘‘error term’’ trees in Lemma 4.1).

As in the odd case, all algebraically canceling pairs of intersections with non- ∞ trees can be exchanged for geometrically canceling pairs by [38]. To finish building the desired order $2n + 1$ twisted Whitney tower, we will describe how to eliminate the remaining algebraically canceling pairs of clean twisted order n Whitney disks (all having simple trees) using a construction that bands together Whitney disks and is additive on twistings. This construction is an iterated elaboration of a construction originally from Chapter 10.8 of [19] (which was used to show that that $\tau_1 \otimes \mathbb{Z}_2$ did not depend on choices of pairing intersections by Whitney disks).




Figure 20: The simple twisted tree J_n^∞ .

Consider an algebraically canceling pair of clean ± 1 -twisted Whitney disks W_{J_n} and W'_{J_n} , whose simple ∞ -trees $+J_n^\infty$ and $-J_n^\infty$ are as in Figure 20, using the notation $J_n = (\cdots((i, j_1), j_2), \cdots, j_n)$. Each trivalent vertex corresponds to a Whitney disk, and we will work from left to right, starting with the order one Whitney disks $W_{(i, j_1)}$ and $W'_{(i, j_1)}$, banding together Whitney disks of the same order from the two trees, while only creating new unpaired intersections of order greater than $2n$. At the last step, W_{J_n} and W'_{J_n} will be banded together into a single framed clean Whitney disk, providing the desired geometric cancellation. (The reason for working with *simple* trees is that the construction for achieving geometric cancellation requires *connected* surfaces for certain steps. For instance, the following construction only gets started because the left most trivalent vertices of an algebraically canceling pair of simple trees correspond to Whitney disks which pair the connected order zero surfaces D_i and D_{j_1} .)




Figure 21: The Whitney disks $W_{(i, j_1)}$, $W'_{(i, j_1)}$, and V' are banded together to form the Whitney disk $W''_{(i, j_1)}$ pairing the outermost pair of intersections between D_i and D_{j_1} . In the cases $n > 1$, the interior of $W''_{(i, j_1)}$ contains two pairs of canceling intersections with D_{j_2} (which are not shown), and supports the sub-towers consisting of the rest of the higher-order Whitney disks (that were supported by $W_{(i, j_1)}$ and $W'_{(i, j_1)}$) corresponding to the trivalent vertices in both trees $\pm J_n^\infty$.

To start the construction consider the Whitney disks $W_{(i, j_1)}$ and $W'_{(i, j_1)}$, pairing intersections between the order zero immersed disks D_i and D_{j_1} . Let V be another Whitney disk for a canceling pair consisting of one point from each of the points paired by $W_{(i, j_1)}$ and $W'_{(i, j_1)}$. Figure 21 illustrates how a parallel copy V' of V can be banded together with $W_{(i, j_1)}$ and $W'_{(i, j_1)}$ to form a Whitney disk $W''_{(i, j_1)}$ for the remaining canceling pair. The twisting of $W''_{(i, j_1)}$ is the sum of the twistings on $W_{(i, j_1)}$, $W'_{(i, j_1)}$, and V ; so $W''_{(i, j_1)}$ is framed if V is framed, since both $W_{(i, j_1)}$ and $W'_{(i, j_1)}$ are framed for $n > 1$ (and in the

$n = 1$ case $W_{(i,j_1)} = W_{J_n}$ and $W'_{(i,j_1)} = W'_{J_n}$ contribute canceling ± 1 twistings). If V is both framed and clean, then the result of replacing $W_{(i,j_1)}$ and $W'_{(i,j_1)}$ by V and $W''_{(i,j_1)}$ preserves the order of \mathcal{W} and creates no new intersections.

So if $n = 1$, then W_{J_n} and W'_{J_n} have been geometrically canceled, meaning that their corresponding ∞ -trees have been eliminated from $t(\mathcal{W})$ without creating any new unpaired order $2n$ intersections or new twisted order n Whitney disks.

The next step shows how V can be arranged to be framed and clean, at the cost of only creating intersections of order greater than $2n$: Any twisting $\omega(V)$ can be killed by boundary twisting V into D_{j_1} . Then, using the construction shown in Figure 22, any interior intersection between V and any K -sheet (e.g. an intersection with D_{j_1} from boundary-twisting) can be pushed down into D_i and paired by a thin Whitney disk $W_{(K,i)}$, which in turn has intersections with the D_{j_1} -sheet that can be paired by a Whitney disk $W_{K_1} := W_{((K,i),j_1)}$ made from a Whitney-parallel copy of $W_{(i,j_1)}$. Now, parallel copies of the Whitney disks from the sub-tower supported by $W_{(i,j_1)}$ can be used to build a sub-tower on W_{K_1} : Using the notation $K_{r+1} = (K_r, i)$, for $r = 1, 2, 3, \dots, n$, the Whitney disk $W_{K_{r+1}}$ is built from a Whitney-parallel copy of W_{J_r} , and pairs intersections between W_{K_r} and j_r . Note that the order of each W_{K_r} is at least r . The top order $W_{K_{n+1}}$ inherits the ± 1 -twisting from W_{J_n} , and has a single interior intersection with tree $\langle K_{n+1}, J_n \rangle$ which is of order at least $2n + 1$.




Figure 22

For multiple intersections between V and various K -sheets this part of the construction can be carried out simultaneously using nested parallel copies of the thin Whitney disk in Figure 22 and more Whitney-parallel copies of the sub-towers described in the previous paragraph.

The result of the construction so far is that the left-most trivalent vertices of the trees $+J_n^\infty$ and $-J_n^\infty$ now correspond to the *same* order 1 Whitney disk $W_{j_1}'' = W_{(i,j_1)}''$, at the cost of having created (after splitting-off) a clean twisted Whitney disk of order at least $n + 1$, and an unpaired intersection of order at least $2n + 1$. In particular, this completes the proof for the case $n = 1$.

For the cases $n > 1$, observe that since W_{j_1}'' is *connected*, this construction can be repeated, with W_{j_1}'' playing the role of D_i , and D_{j_2} playing the role of D_{j_1} , to get a single order 2 Whitney disk W_{j_2}'' which corresponds to the second trivalent vertices from the left in both trees $+J_n^\infty$ and $-J_n^\infty$. By iterating the construction, eventually we band together W_{j_n} and W_{j_n}' into a single framed clean Whitney disk at the last step, having created only clean twisted Whitney disks of order at least $n + 1$, and unpaired intersections of order at least $2n + 1$.

4.2 The geometric twisted IHX relation

The proof of Theorem 2.10 is completed by the following lemma which describes a twisted IHX construction on the intersection tree $t(\mathcal{W})$ of a twisted Whitney tower \mathcal{W} . This geometric move is based on the framed version given in Lemma 7.2 of [34].

Lemma 4.1 *Any split twisted Whitney tower \mathcal{W} containing a clean $+1$ -twisted Whitney disk with signed ∞ -tree $+I^\infty$ can be modified (in a neighborhood of the Whitney disks and local order zero sheets corresponding to I^∞) to a twisted Whitney tower \mathcal{W}' such that $t(\mathcal{W}')$ differs from $t(\mathcal{W})$ exactly by replacing $+I^\infty$ with the signed trees $+H^\infty$, $+X^\infty$, and $-\langle H, X \rangle$, where $+I - H + X$ is a Jacobi relator.*

Similarly, a clean -1 -twisted Whitney disk with ∞ -tree $-I^\infty$ in $t(\mathcal{W})$ can be replaced by $-H^\infty$, $-X^\infty$, and $+\langle H, X \rangle$ in $t(\mathcal{W}')$.

Proof Before describing how to adapt the construction and notation of [34] to give a detailed proof of Lemma 4.1, we explain why the framed geometric relation $+I = +H - X$ leads to the twisted relation $+I^\infty = +H^\infty + X^\infty - \langle H, X \rangle$. In the framed case, a Whitney disk with tree I is replaced by Whitney disks with trees H and X , such that the new Whitney disks are parallel copies of the original using the Whitney framing, and inherit the framing of the original. In order to preserve the trivalent vertex orientations of the trees, the orientation of the H-Whitney disk is the same as the original I-Whitney disk, and the orientation of the X-Whitney disk is the opposite of the I-Whitney disk. Now, if the original I-Whitney disk was $+1$ -twisted, then both the H- and X-Whitney disks will inherit this same $+1$ -twisting, because the twisting – which is

a self-intersection number – is independent of the Whitney disk orientation. The H- and X-Whitney disks will also intersect in a single point with sign -1 , since they inherited opposite orientations from the I-Whitney disk. Thus, (after splitting) a twisted Whitney tower can be modified so that a $+I^\infty$ is replaced by exactly $+H^\infty + X^\infty - \langle H, X \rangle$ in the intersection forest. Similarly, a $-I^\infty$ can be replaced exactly by $-H^\infty - X^\infty + \langle H, X \rangle$.

The framed IHX Whitney move construction is described in detail in [34] (over four pages, including six figures). We describe here how to adapt that construction to the present twisted case, including the relevant modification of notation. Orientation details are not given in [34], but all that needs to be checked is that the X-Whitney disk inherits the opposite orientation as the H-Whitney disk (given that the tree orientations are preserved, and using our negative-corner orientation convention in 2.3 above). In Lemma 7.2 of [34], the “split sub-tower \mathcal{W}_p ” refers to the Whitney disks and order zero sheets containing the tree t_p of an unpaired intersection p in a split Whitney tower \mathcal{W} . In the current setting, a clean $+1$ -twisted Whitney disk W plays the role of p , and the construction will modify \mathcal{W} in a neighborhood of the Whitney disks and order zero sheets containing the ∞ -tree associated to W . In the notation of Figure 18 of [34], the sub-tree of the I-tree denoted by L contains p , so to interpret the entire construction in our case only requires the understanding that this sub-tree contains the ∞ -label sitting in W . (Note that in Figure 18 of [34] the labels I, J, K and L denote *sub-trees*, and in particular the I -labeled sub-tree should not be confused with the “I-tree” in the IHX relation.)

In the case where the L -labeled sub-tree is order zero, then L is just the ∞ -label, and the upper trivalent vertex of the I-tree in Figure 18 of [34] corresponds to the clean $+1$ -twisted W , with ∞ -tree $((I, J), K)^\infty$. Then the construction, 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, yields the $+1$ -twisted H- and X-Whitney disks as discussed in the first paragraph of this proof, with ∞ -trees $(I, (J, K))^\infty$ and $(J, (I, K))^\infty$, and non- ∞ tree $\langle (I, (J, K)), (J, (I, K)) \rangle$ corresponding to the resulting unpaired intersection (created by taking Whitney-parallel copies of the twisted W to form the H- and X-Whitney disks).

In the case where the L -labeled sub-tree is order 1 or greater, then the upper trivalent vertex of the I-tree in Figure 18 of [34] corresponds to a framed Whitney disk, and Whitney-parallel copies of this framed Whitney disk and the other Whitney disks corresponding to the L -labeled sub-tree are also used to construct the sub-towers containing the $+1$ -twisted Whitney disks with H and X ∞ -trees (which will again will lead to a single unpaired intersection as before). \square

4.3 Twisted even order and framed odd order Whitney towers

The following lemma implies that $\mathbb{W}_{2n}^{\text{co}} \subset \mathbb{W}_{2n-1}$, a fact that will be used later in the proof of Theorem 5.1:

Lemma 4.2 *If a collection A of properly immersed surfaces in a simply connected 4 -manifold supports an order $2n$ twisted Whitney tower, then A is homotopic (rel ∂) to A' which supports an order $2n - 1$ framed Whitney tower.*

Proof Let \mathcal{W} be any order $2n$ twisted Whitney tower \mathcal{W} supported by A . If \mathcal{W} contains no order n non-trivially twisted Whitney disks, then \mathcal{W} is an order $2n$ framed Whitney tower, hence also is an order $2n - 1$ framed Whitney tower. If \mathcal{W} does contain order n non-trivially twisted Whitney disks, they can be eliminated at the cost of only creating intersections of order at least $2n - 1$ as follows:

Consider an order n twisted Whitney disk $W_J \subset \mathcal{W}$ with twisting $\omega(W_J) = k \in \mathbb{Z}$. If W_J pairs intersections between an order zero surface A_i and an order $n - 1$ Whitney disk W_I then $J = (i, I)$, and by performing $|k|$ boundary-twists of W_J into W_I , W_J can be made to be framed at the cost of only creating $|k|$ order $2n - 1$ intersections, whose corresponding trees are of the form $\langle (i, I), I \rangle$.

If W_J pairs intersections between two Whitney disks, then by applying the twisted geometric IHX move of Lemma 4.1 (as many times as needed), W_J can be replaced by (many) order n twisted Whitney disks each having a boundary arc on an order zero surface as in the previous case, at the cost of only creating unpaired intersections of order $2n$, each of which is an error term in Lemma 4.1. \square

4.4 Obstruction theory for the reduced tree groups

Using the twisted IHX Lemma 4.1, this section strengthens the obstruction theory for framed Whitney towers described in [38] by showing that the vanishing of $\tau_{2n-1}(\mathcal{W})$ in the reduced group $\tilde{\mathcal{T}}_{2n-1} := \mathcal{T}_{2n-1} / \text{Im}(\Delta_{2n-1})$ is sufficient for the promotion of \mathcal{W} to a Whitney tower of order $2n$. This means that \mathcal{T}_n can be replaced everywhere by $\tilde{\mathcal{T}}_n$ (with $\tilde{\mathcal{T}}_{2n} := \mathcal{T}_{2n}$) throughout Section 3, showing that the realization maps $\tilde{R}_n : \tilde{\mathcal{T}}_n \rightarrow \mathbb{W}_n$ are well-defined epimorphisms.

Recall from the introduction the *framing relations* which in odd orders give the reduced group $\tilde{\mathcal{T}}_{2n-1} := \mathcal{T}_{2n-1} / \text{Im}(\Delta_{2n-1})$ with Δ_{2n-1} given as follows.

Definition 4.3 The map $\Delta_{2n-1} : \mathcal{T}_{n-1} \rightarrow \mathcal{T}_{2n-1}$ is defined for generators $t \in \mathcal{T}_{n-1}$ by

$$\Delta(t) := \sum_{v \in t} \langle \ell(v), (T_v(t), T_v(t)) \rangle$$

where the sum is over all univalent vertices v of t , with $T_v(t)$ denoting the rooted tree gotten by replacing v with a root, and $\ell(v)$ the original label of v .

That Δ_{2n-1} is a well-defined homomorphism is clear since AS and IHX relations go to “doubled” relations. See Figure 23 for explicit illustrations of Δ_1 and Δ_3 .




Figure 23: The map $\Delta_{2n-1} : \mathcal{T}_{n-1} \rightarrow \mathcal{T}_{2n-1}$ in the cases $n = 1$ and $n = 2$.

The following theorem strengthens Theorem 2.6 in Section 2.

Theorem 4.4 *If a collection A of properly immersed surfaces in a simply connected 4-manifold supports a framed Whitney tower \mathcal{W} of order $(2n - 1)$ with $\tau_{2n-1}(\mathcal{W}) \in \text{Im}(\Delta_{2n-1})$, then A is regularly homotopic (rel ∂) to A' which supports a framed Whitney tower of order $2n$.*

Proof As discussed above in the outline the proof of Theorem 2.10 (section 4.1), to prove Theorem 4.4 it will suffice to show that the intersection forest $t(\mathcal{W})$ can be changed by trees representing any element in $\text{Im}(\Delta_{2n-1}) < \mathcal{T}_{2n-1}$ at the cost of only introducing trees of order greater than or equal to $2n$, so that the order $2n - 1$ trees in $t(\mathcal{W})$ all occur in algebraically canceling pairs. Note that $\text{Im}(\Delta_{2n-1})$ is 2-torsion by the AS relations, so orientations and signs are not an issue here. As in Section 4, elements of $t(\mathcal{W})$ will be denoted by formal sums, and \mathcal{W} will not be renamed as modifications are made.

The case $n = 1$: Given any order zero tree $\langle i, j \rangle$, create a clean framed Whitney disk $W_{(i,j)}$ by performing a finger move between the order zero surfaces A_i and A_j . Then

use a twisted finger move (Figure 13) to split $W_{(i,j)}$ into two twisted Whitney disks with associated trees $(i,j)^\infty - (i,j)^\infty$. Now boundary-twist each Whitney disk into a different sheet to recover the framing and add

$$\langle i, (i,j) \rangle + \langle j, (i,j) \rangle = \Delta_1(\langle i,j \rangle)$$

to $t(\mathcal{W})$. Alternatively, after creating the framed $W_{(i,j)}$, perform an interior twist on $W_{(i,j)}$ to get $\omega(W_{(i,j)}) = \pm 2$, then kill $\omega(W_{(i,j)})$ by two boundary-twists, one into each sheet, again adding $\langle i, (i,j) \rangle + \langle j, (i,j) \rangle$ to $t(\mathcal{W})$. Note that $\text{Im } \Delta_1$ in \mathcal{T}_1 corresponds to the order 1 FR framing relation of [36, 37].




Figure 24: Multiple ∞ -roots attached to a tree represent sums (disjoint unions) of trees. On the left: the two trees that result from twist-splitting a clean $W_{(i,(I_1,I_2))}$ in the case $\langle i, (I_1, I_2) \rangle = \langle i, ((j,k), (a, (b,c))) \rangle$. Each arrow indicates an application of a twisted IHX Whitney move, which pushes ∞ -roots towards the univalent vertices. The right-most sum of trees becomes the image of $\langle i, (I_1, I_2) \rangle$ under Δ after applying boundary-twists to the associated twisted Whitney disks.

The cases $n > 1$: For any order $n-1$ tree $\langle i, (I_1, I_2) \rangle$, create a clean $W_{(i,(I_1,I_2))}$ by finger moves. (Here we are taking any order $n-1$ tree, choosing an i -labeled univalent vertex, and writing it as the inner product of the order zero rooted tree i and the remaining order $n-1$ tree.) Then split $W_{(i,(I_1,I_2))}$ using a twisted finger move to get two twisted Whitney disks each having associated ∞ -tree $(i, (I_1, I_2))^\infty$. Leave one of these twisted Whitney disks alone, and to the other apply the twisted geometric IHX Whitney move (Lemma 4.1 of Section 4) to replace $(i, (I_1, I_2))^\infty$ by $(I_1, (I_2, i))^\infty + (I_2, (i, I_1))^\infty - \langle (I_1, (I_2, i)), (I_2, (i, I_1)) \rangle$ in $t(\mathcal{W})$. Note that the tree $\langle (I_1, (I_2, i)), (I_2, (i, I_1)) \rangle$ is order $2n$. If I_1 and I_2 are not both order zero then continue to apply the twisted geometric IHX Whitney move (pushing the ∞ -labeled vertices away from the ∞ -labeled vertex that is adjacent to the original i -labeled vertex) until the resulting union of trees has all ∞ -labeled vertices adjacent to a univalent vertex (all twisted Whitney disks have a boundary arc on an order zero surface) – see Figure 24 for an example. Then, boundary-twisting each twisted Whitney disk into the order zero surface recovers the framing on each Whitney disk and the resulting change in $t(\mathcal{W})$ is a sum of trees as in the right hand side of the equation in Definition 4.3 representing the image of $\langle i, (I_1, I_2) \rangle$ under Δ_{2n-1} , together with trees of order at least $2n$. \square

5 From the twisted to the framed classification

This section fills in the outline of section 1.3 by explaining how to derive the classification of the framed Whitney tower filtration from the classification of the twisted filtration described in section 1.2 of the introduction. The main tool is Theorem 5.1, which relates the two filtrations and their relevant tree groups in a diagram of exact sequences. The subsequent sections 5.1 and 5.2 show how this result, together with the Levine Conjecture and the Milnor invariant-intersection invariant relationship, leads to the framed classification in terms of Milnor invariants, higher-order Sato-Levine invariants and higher-order Arf invariants. The proof of Theorem 5.1 is completed in section 5.4.

The starting point is the following surprisingly simple relation between the twisted and framed Whitney tower filtrations. Recall that in even orders the reduced groups $\tilde{\mathcal{T}}_{2k}$ and realization maps \tilde{R}_{2k} are by definition equal to \mathcal{T}_{2k} and R_{2k} .

Theorem 5.1 *There are commutative diagrams of exact sequences*

$$\begin{array}{ccccccccc}
 0 & \longrightarrow & \tilde{\mathcal{T}}_{2k} & \longrightarrow & \mathcal{T}_{2k}^{\infty} & \longrightarrow & \tilde{\mathcal{T}}_{2k-1} & \longrightarrow & \mathcal{T}_{2k-1}^{\infty} & \longrightarrow & 0 \\
 & & \downarrow \tilde{R}_{2k} & & \downarrow R_{2k}^{\infty} & & \downarrow \tilde{R}_{2k-1} & & \downarrow R_{2k-1}^{\infty} & & \\
 0 & \longrightarrow & \mathcal{W}_{2k} & \longrightarrow & \mathcal{W}_{2k}^{\infty} & \longrightarrow & \mathcal{W}_{2k-1} & \longrightarrow & \mathcal{W}_{2k-1}^{\infty} & \longrightarrow & 0
 \end{array}$$

where all maps in the bottom row are induced by the identity on the set of links. Moreover, there are isomorphisms

$$\text{Cok}(\mathcal{T}_{2k} \rightarrow \mathcal{T}_{2k}^{\infty}) \cong \mathbb{Z}_2 \otimes \mathcal{L}'_{k+1} \cong \text{Ker}(\tilde{\mathcal{T}}_{2k-1} \rightarrow \mathcal{T}_{2k-1}^{\infty})$$

Here $\mathcal{L}' = \bigoplus \mathcal{L}'_n$ is Levine's *quasi-Lie algebra*, which we define using the usual identification of brackets with rooted trees:

Definition 5.2 ([29]) The *degree n* abelian group $\mathcal{L}'_n = \mathcal{L}'_n(m)$ is generated by rooted trees of order $(n - 1)$, each having an unlabeled root univalent vertex, with all other univalent vertices labeled by elements of $\{1, \dots, m\}$, modulo the AS and IHX relations of Figure 5.

The prefix 'quasi' reflects the fact that, although the IHX relation corresponds to the Jacobi identity, the usual Lie algebra self-annihilation relation $[X, X] = 0$ does *not* hold in \mathcal{L}' . It is replaced by the weaker antisymmetry (AS) relation $[Y, X] = -[X, Y]$.

Levine shows that in odd *degrees* (even orders) the natural projection $L'_{2k-1} \rightarrow L_{2k-1}$ is an isomorphism, while in even *degrees* (odd orders) we have the split exact sequence

$$(L) \quad 0 \rightarrow \mathbb{Z}_2 \otimes L_k \rightarrow L'_{2k} \rightarrow L_{2k} \rightarrow 0$$

where the left map sends X to $[X, X]$, and the right map is the natural projection [30].

To see exactness of the bottom row of Theorem 5.1 one first observes that there is a natural inclusion $W_n \subseteq W_n^\infty$, and that $W_{2k-1}^\infty = W_{2k-1}$ by definition. Then, from the inclusion $W_{2k}^\infty \subseteq W_{2k-1}$ shown above in Lemma 4.2, exactness follows since $W_n := W_n/W_{n+1}$ and $W_n^\infty := W_n^\infty/W_{n+1}^\infty$. Proof of Theorem 5.1 is completed in Section 5.4, using the *universality* of \mathcal{T}_{2k}^∞ as the target of quadratic refinements of the canonical ‘inner product’ pairing

$$\langle \cdot, \cdot \rangle : L'_{k+1} \times L'_{k+1} \rightarrow \mathcal{T}_{2k}$$

given by gluing the roots of two rooted trees, see Definition 2.1 and [15].

In order to proceed with the analysis of the framed filtration, notice that the group

$$K_{2k-1}^\mu := \text{Ker}(W_{2k-1} \rightarrow W_{2k-1}^\infty \cong D_{2k-1})$$

is precisely the kernel of the Milnor invariant $\mu_{2k-1} : W_{2k-1} \rightarrow D_{2k-1}$, where the *order n Milnor invariant* $\mu_n : W_n \rightarrow D_n$ is defined in the framed setting via the composition $W_n \rightarrow W_n^\infty \rightarrow D_n$ induced by the inclusion $W_n \subset W_n^\infty$. This breaks the bottom exact sequence in Theorem 5.1 into two short exact sequences, and calculating K_{2k-1}^μ will thus allow us to compute W_{2k} and W_{2k-1} in terms of W_{2k}^∞ and W_{2k-1}^∞ .

On the other hand, the group $K_{2k-1}^\mu \cong W_{2k}^\infty/W_{2k}$ precisely measures the obstructions to framing a twisted Whitney tower of order $2k$. We will show next how these obstructions are detected by the higher-order Sato-Levine invariants defined as projections of the order $2k$ Milnor invariants, together with a direct analogue of the higher-order Arf invariants defined for the framed Whitney tower filtration.

In direct analogy with D_n , the group D'_n is defined as the kernel of the quasi-Lie bracketing map $L'_1 \otimes L'_{n+1} \rightarrow L'_{n+2}$. Analogously to η_n , there is a map $\eta'_n : \mathcal{T}_n \rightarrow D'_n$ defined again by summing over choosing a root at each univalent vertex of each generator:

Definition 5.3 ([29]) On trees $t \in \mathcal{T}_n$ set $\eta'_n(t) := \sum_{v \in t} X_{\ell(v)} \otimes T'_v(t)$, where the sum is over all univalent vertices v of t , with $T'_v(t) \in L'_{n+1}$ denoting the rooted tree gotten by replacing v with a root, and $\ell(v)$ the original label of v .

The previously-mentioned Levine Conjecture [29] is the statement that η'_n is an isomorphism, a surprisingly difficult fact which we prove in [12]:

Theorem 5.4 ([12]) *The maps $\eta'_n: \mathcal{T}_n \rightarrow D'_n$ are isomorphisms.*

The surjectivity of the realization maps (Theorem 1.3) together with the factorization $\eta'_n = \mu_n \circ R_n$ from [11, 38] then gives us the following commutative diagram:

$$\begin{array}{ccc} \mathcal{T}_{2k} & \xrightarrow{R_{2k}} & W_{2k} \\ & \searrow \eta'_{2k} & \downarrow \mu_{2k} \\ & & D'_{2k} \end{array}$$

The groups D'_{2k} are finite index subgroups of D_{2k} and hence are free abelian of known rank for all k [30, Cor.2.3], so by Theorem 5.4 this diagram gives the classification of the Whitney tower filtration in all *even* orders, implying Theorem 1.4:

Theorem 5.5 *The maps R_{2k} and $\mu_{2k}: W_{2k} \rightarrow D'_{2k}$ are isomorphisms.*

5.1 Higher-order Sato-Levine invariants

In order, to define the *order $2k-1$ Sato-Levine invariant* $SL_{2k-1}: K_{2k-1}^\mu \rightarrow \mathbb{Z}_2 \otimes L_{k+1}$, we first use a result of Levine to define homomorphisms sl_{2k} algebraically:

Definition 5.6 The epimorphisms $sl_{2k}: D_{2k} \rightarrow \mathbb{Z}_2 \otimes L_{k+1}$ are defined by the snake lemma applied to the diagram:

$$\begin{array}{ccccc} & & & & \mathbb{Z}_2 \otimes L_{k+1} \\ & & & & \downarrow sq \\ D'_{2k} & \longrightarrow & L_1 \otimes L'_{2k+1} & \longrightarrow & L'_{2k+2} \\ \downarrow & & \downarrow \cong & & \downarrow \\ D_{2k} & \longrightarrow & L_1 \otimes L_{2k+1} & \twoheadrightarrow & L_{2k+2} \\ \downarrow sl_{2k} & & & & \\ \mathbb{Z}_2 \otimes L_{k+1} & & & & \end{array}$$

The two horizontal sequences are exact by definition and the vertical sequence on the right is exact by Theorem 2.2 of [30]. The squaring map on the upper right is $sq(1 \otimes X) := [X, X]$.

Now suppose $L \in \mathbb{W}_{2k-1}$ represents an element in

$$\mathbb{K}_{2k-1}^\mu := \text{Ker}(\mu_{2k-1} : \mathbb{W}_{2k-1} \rightarrow \mathbb{D}_{2k-1}).$$

Since $\mu_{2k-1}(L) = 0$, we have that $\mu_{2k}(L) \in \mathbb{D}_{2k}$ is defined. This gives the definition of the higher-order Sato-Levine invariants:

Definition 5.7 The Sato-Levine invariant $\text{SL}_{2k-1}(L)$ is equal to $s\ell_{2k} \circ \mu_{2k}(L)$.

Summarizing the previous discussion, the cokernel of the inclusion $\mathbb{D}'_{2k} \rightarrow \mathbb{D}_{2k}$ is isomorphic to $\mathbb{Z}_2 \otimes \mathbb{L}_{k+1}$, and we get the following commutative diagram:

$$\begin{array}{ccccc} \mathbb{W}_{2k} & \hookrightarrow & \mathbb{W}_{2k}^\infty & \twoheadrightarrow & \mathbb{K}_{2k-1}^\mu \\ \cong \downarrow \mu_{2k} & & \downarrow \mu_{2k} & & \downarrow \text{SL}_{2k-1} \\ \mathbb{D}'_{2k} & \hookrightarrow & \mathbb{D}_{2k} & \twoheadrightarrow & \mathbb{Z}_2 \otimes \mathbb{L}_{k+1} \end{array}$$

We know from Theorem 1.13 that $\mu_{2k} : \mathbb{W}_{2k}^\infty \rightarrow \mathbb{D}_{2k}$ is an isomorphism when k is even, implying the following proposition:

Proposition 5.8 SL_{4k-1} gives an isomorphism $\mathbb{K}_{4k-1}^\mu \cong \mathbb{Z}_2 \otimes \mathbb{L}_{2k+1}$.

Then, using the fact that $\mathbb{L}_{2k+1} = \mathbb{L}'_{2k+1}$ we get the following commutative diagram from Theorem 5.1 and Theorem 1.13:

$$\begin{array}{ccccc} \mathbb{Z}_2 \otimes \mathbb{L}_{2k+1} & \hookrightarrow & \tilde{\mathcal{T}}_{4k-1} & \twoheadrightarrow & \mathcal{T}_{4k-1}^\infty \\ \downarrow \cong & & \downarrow \tilde{R}_{4k-1} & & \downarrow \cong \\ \mathbb{K}_{4k-1}^\mu & \hookrightarrow & \mathbb{W}_{4k-1} & \twoheadrightarrow & \mathbb{W}_{4k-1}^\infty \end{array}$$

Because $\mathbb{K}_{4k-1}^\mu \cong \mathbb{Z}_2 \otimes \mathbb{L}_{2k+1}$, the left-hand epimorphism of finite dimensional \mathbb{Z}_2 -vector spaces must be an isomorphism. This implies that \tilde{R}_{4k-1} is an isomorphism, and in combination with Theorem 5.5 gives the classification of the framed filtration in three quarters of the cases:

Theorem 5.9 If $n \not\equiv 1 \pmod{4}$, then $\tilde{R}_n : \tilde{\mathcal{T}}_n \rightarrow \mathbb{W}_n$ is an isomorphism.

In particular, this result (together with Proposition 5.8) proves Theorem 1.7 from the introduction.

5.2 Higher-order Arf invariants in the framed filtration

Finally we consider the remaining cases W_{4k-3} of the framed filtration, which is where the higher-order Arf invariants reappear. Let K_{4k-3}^{SL} be the kernel of the order $4k-3$ Sato-Levine invariant $\text{SL}_{4k-3}: K_{4k-3}^{\mu} \rightarrow \mathbb{Z}_2 \otimes L_{2k}$. Recall from Definition 1.15 that in the twisted setting the higher-order Arf invariants $\text{Arf}_k: K_{4k-2}^{\infty} \rightarrow (\mathbb{Z}_2 \otimes L_k) / \text{Ker } \alpha_k^{\infty}$ are defined by inverting a surjection $\alpha_k^{\infty}: \mathbb{Z}_2 \otimes L_k \rightarrow K_{4k-2}^{\infty}$ onto the kernel K_{4k-2}^{∞} of the order $4k-2$ invariant $\mu_{4k-2}: W_{4k-2}^{\infty} \rightarrow D_{4k-2}$.

Lemma 5.10 K_{4k-3}^{SL} is canonically isomorphic to K_{4k-2}^{∞} .

Proof This follows from the commutative diagram:

$$\begin{array}{ccccc}
 & & K_{4k-2}^{\infty} & \dashrightarrow & K_{4k-3}^{\text{SL}} \\
 & & \downarrow & & \downarrow \\
 W_{4k-2} & \rightarrow & W_{4k-2}^{\infty} & \twoheadrightarrow & K_{4k-3}^{\mu} \\
 \cong \downarrow & & \downarrow & & \downarrow \\
 D'_{4k-2} & \rightarrow & D_{4k-2} & \twoheadrightarrow & \mathbb{Z}_2 \otimes L_{2k}
 \end{array}$$

(The left-hand isomorphism comes from Theorem 5.5.) □

Thus the higher-order Arf invariants induce maps $\text{Arf}_k: K_{4k-3}^{\text{SL}} \rightarrow \mathbb{Z}_2 \otimes L_k / \text{Ker}(\alpha_k)$, and we get a complete classification for the framed filtration:

Corollary 5.11 *The groups W_n are classified by Milnor invariants μ_n and in addition, Sato-Levine invariants SL_n if n is odd, and finally, Arf invariants Arf_k for $n = 4k - 3$.*

In particular, a link bounds an order n Whitney tower if and only if it has all vanishing Milnor, Sato-Levine and Arf invariants up to order n (Compare Theorem 1.1).

5.3 The Arf invariant conjecture in terms of reduced realization maps

To complete the translation from the twisted setting, as promised in section 1.3, we observe that Conjecture 1.17 is equivalent to the statement that \tilde{R}_{4k-3} is an isomorphism by considering the following commutative diagram:

$$\begin{array}{ccccc}
 \mathbb{Z}_2 \otimes L'_{2k} & \rightarrow & \tilde{\mathcal{T}}_{4k-3} & \twoheadrightarrow & \mathcal{T}_{4k-3}^{\infty} \\
 \downarrow & & \downarrow \tilde{R}_{4k-3} & & \downarrow \cong \\
 K_{4k-3}^{\mu} & \rightarrow & W_{4k-3} & \twoheadrightarrow & W_{4k-3}^{\infty}
 \end{array}$$

Conjecture 1.17 is equivalent to $K_{4k-2}^{\text{SL}} \cong \mathbb{Z}_2 \otimes L_k$, which in turn is equivalent to $K_{4k-3}^{\mu} \cong L'_{2k} \otimes \mathbb{Z}_2$ (not necessarily canonically), by Levine's exact sequence (L) (see just after Definition 5.2 above).

This happens if and only if the left-hand vertical epimorphism is an isomorphism, which happens if and only if \tilde{R}_{4k-3} is an isomorphism.

Thus, Theorem 5.1 well illustrates both the relationships between the various \mathcal{T} - and W -groups, and the implications of the higher-order Arf invariant Conjecture 1.17. In Section 6 this commutative diagram is extended by the relevant η - and μ -maps to include exact sequences of D -groups, giving our *Master Diagrams*.

5.4 Exactness of the tree sequence

Here we complete the proof of Theorem 5.1 by showing the exactness of the top sequence of \mathcal{T} -groups. The exactness of the bottom sequence of W -groups was checked just after the statement, and the realization maps are well-defined surjections by Section 3. Commutativity of the diagram follows from the fact that elements of the W -groups are determined by the value of the corresponding τ -invariant (Corollary 3.3).

So the proof of Theorem 5.1 will be completed by showing that for any $k \in \mathbb{N}$, there are short exact sequences

$$0 \longrightarrow \mathcal{T}_{2k} \longrightarrow \mathcal{T}_{2k}^{\infty} \longrightarrow \mathbb{Z}_2 \otimes L'_{k+1} \longrightarrow 0$$

and

$$0 \longrightarrow \mathbb{Z}_2 \otimes L'_{k+1} \longrightarrow \tilde{\mathcal{T}}_{2k-1} \longrightarrow \mathcal{T}_{2k-1}^{\infty} \longrightarrow 0.$$

The injectivity of the map $\mathcal{T}_{2k} \hookrightarrow \mathcal{T}_{2k}^{\infty}$ is proven in [15] (see also Remark 40 and Corollary 44 of [13]). The cokernel is then spanned by ∞ -trees, with relations coming from the defining relations of $\mathcal{T}_{2k}^{\infty}$, with non- ∞ trees set to 0: $J^{\infty} = (-J)^{\infty}$, $I^{\infty} = H^{\infty} + X^{\infty}$, and $2J^{\infty} = 0$. Thus the cokernel is isomorphic to $\mathbb{Z}_2 \otimes L'_{k+1}$.

The odd order sequence is shown to be exact as follows:

Recall that the framing relations in $\tilde{\mathcal{T}}_{2k-1}$ are the image of $\Delta_{2k-1}: \mathcal{T}_{k-1} \rightarrow \mathcal{T}_{2k-1}$, as described in Section 4.4. The image of Δ_{2k-1} is 2-torsion by AS relations and hence it factors through $\mathbb{Z}_2 \otimes \mathcal{T}_{k-1}$. Thus we get an exact sequence as in the top of the diagram

below, the middle exact sequence is Corollary 2.3 of [30].

$$\begin{array}{ccccc}
(\mathbb{Z}_2 \otimes \mathcal{T}_{k-1}) / \text{Ker } \Delta & \xrightarrow{\Delta} & \mathcal{T}_{2k-1} & \twoheadrightarrow & \widetilde{\mathcal{T}}_{2k-1} \\
\downarrow & & \cong \downarrow \eta' & & \downarrow \\
\mathbb{Z}_2^m \otimes \mathbb{L}_k & \xrightarrow{sq} & \mathbb{D}'_{2k-1} & \twoheadrightarrow & \mathbb{D}_{2k-1} \\
\parallel & & \downarrow & & \downarrow \\
\mathbb{Z}_2 \otimes \mathbb{L}_1 \otimes \mathbb{L}_k & \xrightarrow{sq} & \mathbb{L}_1 \otimes \mathbb{L}'_{2k} & \twoheadrightarrow & \mathbb{L}_1 \otimes \mathbb{L}_{2k}
\end{array}$$

The map on the right is defined via the factorization $\widetilde{\mathcal{T}}_{2k-1} \twoheadrightarrow \mathcal{T}_{2k-1}^\infty \xrightarrow{\eta} \mathbb{D}_{2k-1}$. So by definition the right-hand square commutes, and induces the left-hand vertical map. In fact, we claim that the induced map $\mathbb{Z}_2 \otimes \mathcal{T}_{k-1} \rightarrow \mathbb{Z}_2^m \otimes \mathbb{L}_k$ factors as

$$\mathbb{Z}_2 \otimes \mathcal{T}_{k-1} \xrightarrow{1 \otimes \eta'} \mathbb{Z}_2 \otimes \mathbb{D}'_{k-1} \rightarrow \mathbb{Z}_2^m \otimes \mathbb{L}_k,$$

with the right hand map induced by $\mathbb{D}'_{k-1} \rightarrow \mathbb{Z}^m \otimes \mathbb{L}'_k \twoheadrightarrow \mathbb{Z}^m \otimes \mathbb{L}_k$. To see this, let $t \in \mathcal{T}_{k-1}$ and compute

$$\begin{aligned}
\eta'(\Delta(1 \otimes t)) &= \eta' \left(\sum_{v \in t} \langle \ell(v), (T_v(t), T_v(t)) \rangle \right) \\
&= \sum_{v \in t} X_{\ell(v)} \otimes (T_v(t), T_v(t)) \\
&= \sum_{v \in t} X_{\ell(v)} \otimes sq(T_v(t)) \\
&= sq(1 \otimes \eta')(1 \otimes t)
\end{aligned}$$

Now we claim that for all orders k , there is an exact sequence

$$\mathbb{Z}_2 \otimes \mathbb{D}'_{k-1} \rightarrow \mathbb{Z}_2^m \otimes \mathbb{L}_k \rightarrow \mathbb{Z}_2 \otimes \mathbb{L}'_{k+1} \rightarrow 0.$$

This is clear if k is odd, by tensoring the defining exact sequence for \mathbb{D}'_{k-1} with \mathbb{Z}_2 . If k is even, then this follows since there is a surjection $\mathbb{D}'_{k-1} \twoheadrightarrow \mathbb{D}_{k-1}$ and $\mathbb{L}'_{k+1} \cong \mathbb{L}_{k+1}$. Therefore, the commutative diagram above supports a vertical short exact sequence on

the left:

$$\begin{array}{ccccc}
 & & & & \mathbb{Z}_2 \otimes L'_{k+1} \\
 & & & & \downarrow \\
 & & & & \vdots \\
 & & & & \downarrow \\
 (\mathbb{Z}_2 \otimes \mathcal{T}_{k-1}) / \text{Ker } \Delta & \xrightarrow{\Delta} & \mathcal{T}_{2k-1} & \longrightarrow & \tilde{\mathcal{T}}_{2k-1} \\
 \downarrow & & \downarrow \eta' & & \downarrow \\
 \mathbb{Z}_2^m \otimes L_k & \longrightarrow & D'_{2k-1} & \longrightarrow & D_{2k-1} \\
 \downarrow & & \downarrow & & \downarrow \\
 & & \mathbb{Z}_2 \otimes L'_{k+1} & &
 \end{array}$$

which gives us the indicated map on the right. Furthermore, $\eta: \mathcal{T}_{2k-1}^\infty \rightarrow D_{2k-1}$ is an isomorphism by Theorem 6.5 below, so the right-hand exact sequence is precisely the exact sequence we're interested in.

6 Summary of computations for the Whitney filtrations

In this section the commutative diagram of \mathcal{T} - and W -groups from Theorem 5.1 is extended by the relevant η - and μ -maps to include exact sequences of D -groups (section 6.1). Then section 6.3 establishes some implications of the Levine Conjecture that were used earlier, including a proof of Proposition 1.14 from the introduction.

6.1 The easy Master Diagram

We have already extensively used the commutative diagram of Theorem 5.1 connecting the 4-term exact sequences for the various \mathcal{T} - and W -groups. Here and in the subsequent subsection we introduce an exact sequence of D -groups to complement these. For this we need to define two additional groups, \tilde{D}_{2k-1} and D_{4k-2}^∞ . The group \tilde{D}_{2k-1} is defined to be the quotient of D'_{2k-1} by the image of the framing relations under η'_{2k-1} . The group D_{4k-2}^∞ , explained in Definition 6.3, is (non-canonically) isomorphic to $D_{4k-2} \oplus (\mathbb{Z}_2 \otimes L_k)$.

Then the entire classification picture can be organized into *Master Diagrams* relating the various aspects of the story. There will be two such diagrams, each covering half of the cases, We present first the ‘‘easier’’ diagram, which combines three commutative triangles we have already seen, together with a new one involving \tilde{D}_{4k-1} . This new

triangle follows more-or-less by definition: since \tilde{R}_{4k-1} is an isomorphism, we can let $\tilde{\mu}_{4k-1} := \tilde{\eta}_{4k-1} \circ \tilde{R}_{4k-1}^{-1}$. The maps $D'_{4k} \rightarrow D_{4k}$ and $\tilde{D}_{4k-1} \rightarrow D_{4k-1}$ are clear, and commutativity follows by definition of the various η maps. The map $D_{4k} \rightarrow \tilde{D}_{4k-1}$ can be defined so that it is induced by the map $\mathcal{T}_{4k}^\infty \rightarrow \tilde{\mathcal{T}}_{4k-1}$.

Theorem 6.1 *The following is a commutative diagram connecting three 4-term exact sequences with triangles of isomorphisms.*

$$\begin{array}{ccccccc}
 \mathcal{T}_{4k} & \longrightarrow & \mathcal{T}_{4k}^\infty & \longrightarrow & \tilde{\mathcal{T}}_{4k-1} & \longrightarrow & \mathcal{T}_{4k-1}^\infty \\
 \downarrow \eta'_{4k} & \searrow \tilde{R}_{4k} & \downarrow & \searrow R_{4k}^\infty & \downarrow & \searrow \tilde{R}_{4k-1} & \downarrow & \searrow R_{4k-1}^\infty \\
 & & W_{4k} & \xrightarrow{\eta_{4k}} & W_{4k}^\infty & \xrightarrow{\tilde{\eta}_{4k-1}} & W_{4k-1} & \xrightarrow{\eta_{4k-1}} & W_{4k-1}^\infty \\
 & \swarrow \mu_{4k} & \downarrow & \swarrow \mu_{4k} & \downarrow & \swarrow \tilde{\mu}_{4k-1} & \downarrow & \swarrow \mu_{4k-1} \\
 D'_{4k} & \longrightarrow & D_{4k} & \longrightarrow & \tilde{D}_{4k-1} & \longrightarrow & D_{4k-1}
 \end{array}$$

Moreover, the three horizontal sequences can each be split into two short exact sequences, with the term $K_{4k-1}^\mu \cong \mathbb{Z}_2 \otimes L_{2k+1}$ appearing in the middle as the cokernel of the left-hand maps and the kernel of the right-hand maps.

Theorem 6.1 will be made precise and proven in the next subsection.

6.2 The hard Master Diagram

The other half of cases are covered by the following diagram. Here we notice that $\mathbb{Z}_2 \otimes L'_{2k}$ breaks the \mathcal{T} -sequence into two exact sequences, but that the exact sequence at the bottom $D'_{4k-2} \rightarrow D_{4k-2} \rightarrow \mathbb{Z}_2 \otimes L_{2k}$ involves L_{2k} and not L'_{2k} . Hence to make the corresponding 4-term D sequence, we create D_{4k-2}^∞ by a pullback diagram as pictured by the ‘‘p.b.-’’-labeled parallelogram. Then one can lift η_{4k-2} to η_{4k-2}^∞ which becomes an isomorphism (Theorem 6.5). We have seen that μ_{4k-2} can be defined with image in D_{4k-2} . Lifting it to μ_{4k-2}^∞ is equivalent to Conjecture 1.17, since that conjecture is equivalent to R_{4k-2}^∞ being an isomorphism. If the lifted maps μ_{4k-2}^∞ exist, they would automatically be isomorphisms.

Theorem 6.2 *The following diagram commutes.*

$$\begin{array}{ccccccc}
 \mathcal{T}_{4k-2} & \longrightarrow & \mathcal{T}_{4k-2}^\infty & \longrightarrow & \tilde{\mathcal{T}}_{4k-3} & \longrightarrow & \mathcal{T}_{4k-3}^\infty \\
 \downarrow \eta_{4k-2} & \searrow \tilde{R}_{4k-2} & \downarrow \eta_{4k-2}^\infty & \searrow R_{4k-2}^\infty & \downarrow \tilde{\eta}_{4k-3} & \searrow \tilde{R}_{4k-3} & \downarrow \eta_{4k-3} \\
 \mathcal{W}_{4k-2} & \longrightarrow & \mathcal{W}_{4k-2}^\infty & \longrightarrow & \mathcal{W}_{4k-3} & \longrightarrow & \mathcal{W}_{4k-3}^\infty \\
 \downarrow \mu_{4k-2} & \swarrow \mu_{4k-2} & \downarrow \mu_{4k-2}^\infty & \swarrow \mu_{4k-2} & \downarrow \tilde{\mu}_{4k-3} & \swarrow \tilde{\mu}_{4k-3} & \downarrow \mu_{4k-3} \\
 D'_{4k-2} & \longrightarrow & D_{4k-2}^\infty & \longrightarrow & \tilde{D}_{4k-3} & \longrightarrow & D_{4k-3} \\
 \parallel & & \downarrow & & \downarrow & & \downarrow \\
 D'_{4k-2} & \longrightarrow & D_{4k-2} & \xrightarrow{\text{p.b.}} & L'_{2k} \otimes \mathbb{Z}_2 & \longrightarrow & L_{2k} \otimes \mathbb{Z}_2
 \end{array}$$

As in the first diagram, the three 4-term horizontal sequences are exact. The horizontal \mathcal{T} - and D -sequences actually break into two short exact sequences with the groups $\mathbb{Z}_2 \otimes L'_{2k}$ in the middle. For the horizontal \mathcal{W} -sequence, the group K_{4k-3}^μ sits in the middle, which we conjecture to be isomorphic to $\mathbb{Z}_2 \otimes L'_{2k}$. At the bottom, there are two short exact sequences, both starting with D'_{4k-2} because the diagonal square is a pullback (p.b.).

In the rest of this subsection we will collect the remaining algebraic definitions and proofs needed to set up these master diagrams.

First we need the definitions of \tilde{D}_{2k-1} and D_{4k-2}^∞ :

Definition 6.3 The groups D_{4k-2}^∞ (as well as the maps $sl'_{4k-2}, p_{2k}, sq^\infty$) are defined by the pullback diagram

$$(D^\infty) \quad \begin{array}{ccccc}
 \mathbb{Z}_2 \otimes L_k & \xrightarrow{sq^\infty} & D_{4k-2}^\infty & \xrightarrow{p_{4k-2}} & D_{4k-2} \\
 \parallel & & \downarrow sl'_{4k-2} & & \downarrow sl_{4k-2} \\
 \mathbb{Z}_2 \otimes L_k & \xrightarrow{sq} & \mathbb{Z}_2 \otimes L'_{2k} & \xrightarrow{p} & \mathbb{Z}_2 \otimes L_{2k}
 \end{array}$$

To define the group \tilde{D}_{2k-1} , we note that the homomorphism $\text{id} \otimes sq: \mathbb{Z}_2^m \otimes L_k \rightarrow \mathbb{Z}_2^m \otimes L_{2k}$ restricts to a homomorphism $\mathbb{Z}_2 \otimes D'_{k-1} \rightarrow D'_{2k-1}$, because elements of the form $X \otimes [Y, Y]$ are in the kernel of the bracketing map. Now \tilde{D}_{2k-1} is defined to be

the quotient of D'_{2k-1} by the image of this homomorphism. It is not hard to show that the image of this homomorphism is equal to the image $\eta'_{2k-1} \circ \Delta(\mathbb{Z}_2 \otimes \mathcal{T}_{k-1})$. Thus there is an induced map $\tilde{\eta}_{2k-1}: \tilde{\mathcal{T}}_{2k-1} \rightarrow \tilde{D}_{2k-1}$.

Lemma 6.4 *There is a canonical lift η_{4k-2}^∞ of η_{4k-2}*

$$\begin{array}{ccc} & & D_{4k-2}^\infty \\ & \nearrow \eta_{4k-2}^\infty & \downarrow p_{4k-2} \\ \mathcal{T}_{4k-2}^\infty & \xrightarrow{\eta_{4k-2}} & D_{4k-2} \end{array}$$

such that $\eta_{4k-2}^\infty((J, J)^\infty) = sq^\infty(1 \otimes J)$ for all rooted trees $J \in \mathcal{L}_k$.

Proof To construct η_{4k-2}^∞ it suffices to observe that we have a commutative diagram

$$\begin{array}{ccc} \mathcal{T}_{4k-2}^\infty & \xrightarrow{\eta_{4k-2}} & D_{4k-2} \\ \downarrow & & \downarrow \\ \mathbb{Z}_2 \otimes \mathcal{L}'_{2k} & \xrightarrow{p} & \mathbb{Z}_2 \otimes \mathcal{L}_{2k} \end{array}$$

which gives rise to a map to the pullback D_{4k-2}^∞ from diagram D^∞ in Definition 6.3 above. To calculate $\eta_{4k-2}^\infty((J, J)^\infty)$ notice that $\eta_{4k-2}((J, J)^\infty) = 0$ and $sl'_{4k-2}((J, J)^\infty) = 1 \otimes (J, J) = sq(1 \otimes J)$. So $\eta_{4k-2}^\infty((J, J)^\infty) = sq^\infty(1 \otimes J)$. \square

Remark 6.1 The superscripts in our η -maps reflect those of their *target* groups.

6.3 Variations on the Levine conjecture

The next theorem establishes that various versions of the η map are isomorphisms as a result of the Levine Conjecture (Theorem 5.4 above), and gives as a corollary the characterization of the kernel of η_{4k-2} stated in Proposition 1.14 of the introduction.

Theorem 6.5 *The following maps are all isomorphisms:*

- (i) $\tilde{\eta}_{2k-1}: \tilde{\mathcal{T}}_{2k-1} \rightarrow \tilde{D}_{2k-1}$
- (ii) $\eta_{2k-1}: \mathcal{T}_{2k-1}^\infty \rightarrow D_{2k-1}$
- (iii) $\eta_{4k}: \mathcal{T}_{4k}^\infty \rightarrow D_{4k}$
- (iv) $\eta_{4k-2}^\infty: \mathcal{T}_{4k-2}^\infty \rightarrow D_{4k-2}^\infty$

Proof (i) follows from Theorem 5.4 and the definition of $\tilde{\eta}_{2k-1}$.

To show (ii), consider the following diagram (commutative by definition of η):

$$\begin{array}{ccccccc}
0 & \longrightarrow & \text{span}\{\langle i, (J, J) \rangle\} & \longrightarrow & \mathcal{T}_{2k-1} & \longrightarrow & \mathcal{T}_{2k-1}^\infty \longrightarrow 0 \\
& & \downarrow \eta' & & \cong \downarrow \eta' & & \downarrow \eta \\
0 & \longrightarrow & \mathbb{Z}_2^m \otimes L_k & \xrightarrow{sq} & D'_{2k-1} & \longrightarrow & D_{2k-1} \longrightarrow 0
\end{array}$$

The bottom row is exact by Corollary 2.3 of [30]. The top row is exact by definition and the middle map η' is an isomorphism by Theorem 5.4. This implies the left-hand restriction η' is one to one, and it is onto since $\eta'(\langle i, (J, J) \rangle) = X_i \otimes J$. Therefore, the right-hand map η is an isomorphism by the 5-lemma.

For (iii), consider the following diagram (commutative by Lemma 6.7)

$$\begin{array}{ccccccc}
0 & \longrightarrow & \mathcal{T}_{4k} & \longrightarrow & \mathcal{T}_{4k}^\infty & \longrightarrow & \mathbb{Z}_2 \otimes L'_{2k+1} \longrightarrow 0 \\
& & \cong \downarrow \eta'_{4k} & & \downarrow \eta_{4k} & & \cong \downarrow p \\
0 & \longrightarrow & D'_{4k} & \longrightarrow & D_{4k} & \xrightarrow{s\ell_{4k}} & \mathbb{Z}_2 \otimes L_{2k+1} \longrightarrow 0
\end{array}$$

The bottom horizontal sequence is exact by Definition 5.6. The top one is part of Theorem 5.1, proven in Section 5.4. Since $L_{2k+1} \cong L'_{2k+1}$ it follows that η_{4k} is an isomorphism.

For (iv), we note that Diagram D^∞ and Lemma 6.4 imply a commutative diagram:

$$\begin{array}{ccccccc}
0 & \longrightarrow & \langle (J, J)^\infty \rangle & \longrightarrow & \mathcal{T}_{4k-2}^\infty & \xrightarrow{\eta_{4k-2}} & D_{4k-2} \longrightarrow 0 \\
& & \downarrow & & \downarrow \eta_{4k-2}^\infty & & \parallel \\
0 & \longrightarrow & \mathbb{Z}_2 \otimes L_k & \xrightarrow{sq_{4k-2}^\infty} & D_{4k-2}^\infty & \xrightarrow{p_{4k-2}} & D_{4k-2} \longrightarrow 0
\end{array}$$

where the vertical left hand map sends $\langle (J, J)^\infty \rangle$ to $1 \otimes J$. The right-hand square commutes by the main commutative triangle of Lemma 6.4, whereas the left square commutes by the calculation $\eta_{4k-2}^\infty(\langle (J, J)^\infty \rangle) = sq_{4k-2}^\infty(1 \otimes J)$, as in Lemma 6.7. Thus η_{4k-2}^∞ is an isomorphism by the 5-lemma. \square

As a corollary to the above proof of (iv) we get the following characterization of $\text{Ker}(\eta_{4k-2}: \mathcal{T}_{4k-2}^\infty \rightarrow D_{4k-2})$ which is equivalent to Proposition 1.14 in the introduction:

Corollary 6.6 *The kernel of the homomorphism $\eta_{4k-2}: \mathcal{T}_{4k-2}^\infty \rightarrow D_{4k-2}$ is isomorphic to $\mathbb{Z}_2 \otimes L_k$, with $\langle (J, J)^\infty \rangle \mapsto 1 \otimes J$.*

Proof That there is an isomorphism follows since η_{4k-2}^∞ is an isomorphism which is a lift of η_{4k-2} and from the exact sequence $0 \rightarrow \mathbb{Z}_2 \otimes L_k \rightarrow D_{4k-2}^\infty \rightarrow D_{4k-2} \rightarrow 0$. The fact that $(J, J)^\infty \mapsto 1 \otimes J$ follows from Lemma 6.4 \square

The following lemma was used in the proof of Theorem 6.5.

Lemma 6.7 *Sending J^∞ to $1 \otimes J$ gives a commutative diagram*

$$\begin{array}{ccc} \mathcal{T}_{2k}^\infty & \longrightarrow & \mathbb{Z}_2 \otimes L'_{k+1} \\ \downarrow \eta_{2k} & & \downarrow p \\ D_{2k} & \xrightarrow{sl_{2k}} & \mathbb{Z}_2 \otimes L_{k+1} \end{array}$$

Proof First, we need a better handle on the map $sl_{2k}: D_{2k} \rightarrow \mathbb{Z}_2 \otimes L_{k+1}$. Let $Z \in D_{2k}$ and pick a lift $Z' \in L_1 \otimes L'_{2k+1}$. Tracing through the snake lemma diagram in Definition 5.6, one sees that the bracket of Z' is a sum of commutators $[J_i, J_i]$, and that $sl_{2k}(Z) = \sum 1 \otimes J_i$.

Consider a tree $t \in \mathcal{T}_{2k}^\infty$ which maps to zero in $\mathbb{Z}_2 \otimes L'_{k+1}$ by definition. Mapping t down by η_{2k} , we end up in D'_{2k} and hence in the kernel of sl_{2k} .

Now consider $J^\infty \in \mathcal{T}_{2k}^\infty$. Then $\eta_{2k}(J^\infty)$ doubles J to $\langle J, J \rangle$ and sums over putting a root at all of the leaves of one copy of J . The result represents an element in $L_1 \otimes L_{k+1}$. Calculating the bracket has the effect of summing over putting a root near all of the leaves on one copy of J in $\langle J, J \rangle$, which modulo IHX is equal to (J, J) . To see this requires pushing the central root of (J, J) to one side using IHX relations.

Thus $sl_{2k}(\eta_{2k}(J^\infty)) = 1 \otimes J$ which is equal to mapping J^∞ right and then down. \square

References

- [1] **J C Cha**, *Link concordance, homology cobordism, and Hirzebruch-type defects from iterated p -covers*, J. Eur. Math. Soc. 12 (2010), 555–610.
- [2] **T Cochran**, *Derivatives of links, Milnor's concordance invariants and Massey products*, Mem. Amer. Math. Soc. Vol. 84 No. 427 (1990).
- [3] **T Cochran**, *k -cobordism for links in S^3* , Trans. Amer. Math. Soc. 327 no 2 (1991) 641–654.
- [4] **T Cochran, S Harvey, C Leidy**, *Knot concordance and higher-order Blanchfield duality*, Geom. Topol. 13 (2009), no. 3, 1419–1482.

- [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 η -invariants*, Duke Math. J. 137 (2007), no. 2, 337–379.
- [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*, Math. Annalen 328 (2004) 135–171.
- [9] **J Conant, R Schneiderman, P Teichner**, *Jacobi identities in low-dimensional topology*, Compositio Mathematica 143 Part 3 (2007) 780–810.
- [10] **J Conant, R Schneiderman, P Teichner**, *Higher-order intersections in low-dimensional topology*, Proc. Natl. Acad. Sci. USA 2011 108 (20) 8081–8084.
- [11] **J Conant, R Schneiderman, P Teichner**, *Milnor Invariants and Twisted Whitney Towers*, preprint (2010), math.GT arXiv.
- [12] **J Conant, R Schneiderman, P Teichner**, *Tree homology and a conjecture of Levine*, To appear in Geom. Topol., preprint (2010) <http://arxiv.org/abs/1012.2780>
- [13] **J Conant, R Schneiderman, P Teichner**, *Universal quadratic refinements and untwisting Whitney towers*, preprint (2010) <http://arxiv.org/abs/1101.3480>
- [14] **J Conant, R Schneiderman, P Teichner**, *Geometric filtrations of string links and homology cylinders*, preprint (2010).
- [15] **J Conant, R Schneiderman, P Teichner**, *Universal quadratic refinements and Whitney towers*, In preparation (2011).
- [16] **M Freedman**, *The disk theorem for four-dimensional manifolds*, Proc. ICM Warsaw (1983) 647–663.
- [17] **M Freedman**, *The topology of four-dimensional manifolds*, J. Diff. Geom. 17 (1982) 357–453.
- [18] **M Freedman, R Kirby**, *A geometric proof of Rochlin’s theorem*, Sympos. Pure Math. Vol. 80, Amer. Math. Soc. (1977) 195–207.
- [19] **M Freedman, F Quinn**, *The topology of 4-manifolds*, Princeton Math. Series 39 Princeton, NJ, (1990).
- [20] **M Freedman, P Teichner**, *4-manifold topology II: Dwyer’s filtration and surgery kernels*, Invent. Math. 122 (1995) 531–557.
- [21] **M Goussarov**, *Finite type invariants and n -equivalence of 3-manifolds*, C. R. Acad. Sci. Paris Ser. I. Math. 329 (1999) 517–522.
- [22] **K Habiro**, *Claspers and finite type invariants of links*, Geom. Topol. 4 (2000) 1–83.
- [23] **N Habegger, G Masbaum**, *The Kontsevich integral and Milnor’s invariants*, Topology 39 (2000) 1253–1289.

- [24] **N Habegger, W Pitsch**, *Tree level Lie algebra structures of perturbative invariants*, J. Knot Theory Ramifications 12 (2003) 333–345.
- [25] **K Igusa, K Orr**, *Links, pictures and the homology of nilpotent groups*, Topology 40 (2001) 1125–1166.
- [26] **V Krushkal**, *Exponential separation in 4-manifolds*, Geometry and Topology, Vol. 4 (2000) 397–405.
- [27] **S Krushkal, P Teichner**, *Alexander duality, Gropes and link homotopy*, Geom. Topol. 1 (1997) 51–69.
- [28] **J Levine**, *Homology cylinders: an enlargement of the mapping class group*, Alg. and Geom. Topology 1 (2001) 243–270.
- [29] **J Levine**, *Addendum and correction to: Homology cylinders: an enlargement of the mapping class group*, Alg. and Geom. Topology 2 (2002) 1197–1204.
- [30] **J Levine**, *Labeled binary planar trees and quasi-Lie algebras*, Alg. and Geom. Topology 6 (2006) 935–948.
- [31] **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.
- [32] **J-B Meilhan, A Yasuhara**, *Characterization of finite type string link invariants of degree < 5* , Math. Proc. Cambridge Philos. Soc. 148 (2010), no. 3, 439–472.
- [33] **K Orr**, *Homotopy invariants of links*, Invent. Math. 95 17 (1989) 379–394.
- [34] **R Schneiderman**, *Whitney towers and Gropes in 4-manifolds*, Trans. Amer. Math. Soc. 358 (2006), 4251–4278.
- [35] **R Schneiderman**, *Simple Whitney towers, half-gropes and the Arf invariant of a knot*, Pacific Journal of Mathematics, Vol. 222, No. 1, Nov (2005).
- [36] **R Schneiderman**, *Stable concordance of knots in 3-manifolds*, Alg. and Geom. Topology 10 (2010) 37–432.
- [37] **R Schneiderman, P Teichner**, *Higher order intersection numbers of 2-spheres in 4-manifolds*, Alg. and Geom. Topology 1 (2001) 1–29.
- [38] **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.
- [39] **A Scorpan**, *The Wild World of 4-Manifolds*, American Mathematical Society (2005).
- [40] **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.
- [41] **P Teichner**, *Knots, von Neumann Signatures, and Grope Cobordism*. Proceedings of the International Congress of Math. Vol II: Invited Lectures (2002) 437–446.
- [42] **P Teichner**, *What is ... a grope?* Notices of the A.M.S. Vol. 54 no. 8 Sep 2004, 894–895.

- [43] **M Yamasaki**, *Whitney's trick for three 2-dimensional homology classes of 4-manifolds*, Proc. Amer. Math. Soc. 75 (1979) 365–371.
- [44] **H Whitney**, *The self intersections of a smooth n -manifold in $2n$ -space*, Annals of Math. 45 (1944) 220–246.

Dept. of Mathematics, University of Tennessee, Knoxville, TN, 37996

Dept. of Mathematics and Computer Science, Lehman College, City University of New York, Bronx, NY 10468

*Dept. of Mathematics, University of California, Berkeley, CA and
Max-Planck Institut für Mathematik, Bonn, Germany*

jconant@math.utk.edu, robert.schneiderman@lehman.cuny.edu,
teichner@mac.com