Journal of Topology 7 (2014) 187–224
doi:10.1112/jtopol/jtt025

Milnor invariants and twisted Whitney towers

J. Conant, R. Schneiderman and P. Teichner

Abstract

This paper describes the relationship between the first non-vanishing *Milnor invariants* of a classical link and the intersection invariant of a *twisted Whitney tower*. This is a certain 2-complex in the 4-ball, built from immersed disks bounded by the given link in the 3-sphere together with finitely many 'layers' of Whitney disks.

The *intersection invariant* is a higher-order generalization of the intersection number between two immersed disks in the 4-ball, well known to give the linking number of the link on the boundary, which measures intersections among the Whitney disks and the disks bounding the given link, together with information that measures the twists (framing obstructions) of the Whitney disks.

This interpretation of Milnor invariants as higher-order intersection invariants plays a key role in our classifications [J. Conant, R. Schneiderman and P. Teichner, 'Higher-order intersections in low-dimensional topology', *Proc. Natl Acad. Sci. USA* 108 (2011) 8131–8138; J. Conant, R. Schneiderman and P. Teichner, 'Whitney tower concordance of classical links', *Geom. Topol.* 16 (2012) 1419–1479] of both the framed and twisted Whitney tower filtrations on link concordance. Here, we show how to realize the *higher-order Arf invariants*, which also play a role in the classifications, and derive new geometric characterizations of links with vanishing length at most $2k$ Milnor invariants.

1. *Introduction*

In his early work [**37**, **38**], Milnor showed how to extract invariants of classical links from nilpotent quotients of the link group. Roughly speaking, Milnor observed that the linking numbers $\mu_L(i,j) \in \mathbb{Z}$ of an oriented link $L = L_1 \cup L_2 \cup \cdots \cup L_m \subset S^3$ vanish if and only if the link group $\pi_1(S^3 \smallsetminus L)$ is isomorphic *modulo 3-fold commutators* to the free group on m generators (the link group of the trivial link). Using this isomorphism, Milnor defined his triple invariants $\mu_L(i,j,k) \in \mathbb{Z}$ which vanish if and only if $\pi_1(S^3 \smallsetminus L)$ is free *modulo 4-fold commutators*. Iterating, he obtained a filtration on the set $\mathbb{L}$ of oriented links in the 3-sphere:

$$\cdots \subseteq \mathbb{M}_3 \subseteq \mathbb{M}_2 \subseteq \mathbb{M}_1 \subseteq \mathbb{L},$$

where a link L lies in $\mathbb{M}_n$ if and only if for all $k \leqslant n$ the Milnor invariants $\mu_L(i_0, i_1, \ldots, i_k) \in \mathbb{Z}$ are defined and vanish. This in turn is equivalent to $\pi_1(S^3 \smallsetminus L)$ being free modulo $(n+2)$-fold commutators, and then the next set of Milnor invariants of *length* $(n+2)$ are defined (via Magnus expansions of the longitudes, thought of as elements in the free group). We refer to Subsection 1.1 for a precise definition and some history on Milnor invariants.

In this paper, we will provide a geometric interpretation of Milnor invariants in terms of the intersection invariants of *twisted Whitney towers*. These are certain 2-complexes in the 4-ball, built on immersed disks bounded by the link L by recursively adding layers of Whitney disks which pair intersections among the previous layers. The intersection invariants measure

Received 24 February 2012; revised 25 April 2013; published online 23 July 2013.

2010 *Mathematics Subject Classification* 57M25 (primary).

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 and a grant (#208938) from the Simons Foundations.

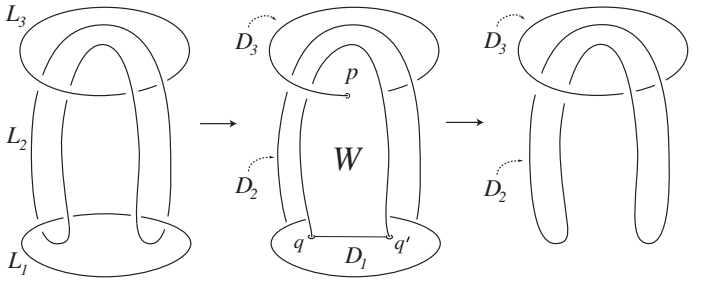

FIGURE 1. *Moving radially into B^4 from left to right, this sequence of pictures shows the Borromean Rings $L = L_1 \cup L_2 \cup L_3 \subset S^3 = \partial B^4$ bounding an order 1 twisted Whitney tower $\mathcal{W} \subset B^4$. It consists of embedded disks D_i with $\partial D_i = L_i$ together with a Whitney disk W that pairs the two intersection points q and q' between D_1 and D_2. The disk D_1 consists of the 'horizontal' opaque disk in the lower part of the middle picture extended by an annular collar back to L_1 in the left picture. Disks D_2 and D_3 consist of the embedded annuli which are the products of L_2 and L_3 with the radial coordinate into B^4 together with embedded disks (not shown) extending further into B^4 that bound the unlink in the right picture. The embedded Whitney disk W is completely contained in the middle picture and has a single intersection point p with D_3. This unpaired 'higher-order' intersection point p shows that $\mu_L(1,2,3) = \pm 1$ and prevents a Whitney move that would promote $\mathcal{W}$ to a collection of slice disks for L.*

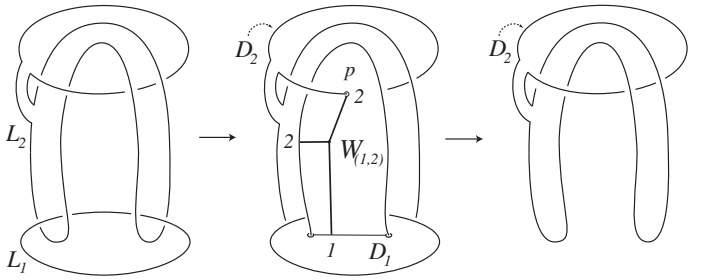

FIGURE 2. *Moving radially into B^4 from left to right, this sequence of pictures shows a Whitehead link $L \subset S^3$ bounding an order 1 Whitney tower $\mathcal{W} \subset B^4$ with $\tau_1(\mathcal{W}) = 1 \multimap \frac{2}{2}$. The left picture shows $L = L_1 \cup L_2 \subset S^3$ formed by an internal band sum on the Borromean rings. Moving into B^4 in the middle and right pictures, L_1 and L_2 bound embedded disks D_1 and D_2 with an embedded Whitney disk $W_{(1,2)}$ pairing $D_1 \cap D_2$. Disk D_1 consists of the 'horizontal' opaque disk in the lower part of the middle picture extended by an annular collar back to L_1 in the left picture. The embedded Whitney disk $W_{(1,2)}$ is completely contained in the middle picture, which also contains the tree $t_p = 1 \multimap \frac{2}{2}$ associated to the unpaired intersection point p between D_2 and the interior of $W_{(1,2)}$. Disk D_2 is the union of the embedded annulus (visible in all three pictures) which is the product of L_2 with the radial coordinate into B^4, together with an embedded disk (not shown) extending further into B^4 that bounds the parallel of L_2 in the right picture.*

'higher-order intersections' among the Whitney disks, as well as certain relative Euler numbers of their normal bundles, and the relationship with the Milnor invariants provides obstructions to 'raising the order' of a Whitney tower (see Figures 1–3 and Section 2). This will show precisely how Milnor invariants are related to the failure of the *Whitney move* [**53**].

As observed in [**28**], the general failure of the Whitney move in smooth 4-manifolds goes back to Rohlin's theorem [**42**]. It was dramatically underlined by Donaldson's restrictions on the intersection form of a smooth 4-manifold [**14**]. Freedman's celebrated recovery [**16**] of the Whitney move in the settings of surgery and the s-cobordism theorem for topological 4-manifolds with 'good' fundamental group was built on a notion of infinitely iterated towers of disks pioneered by Casson [**3**].

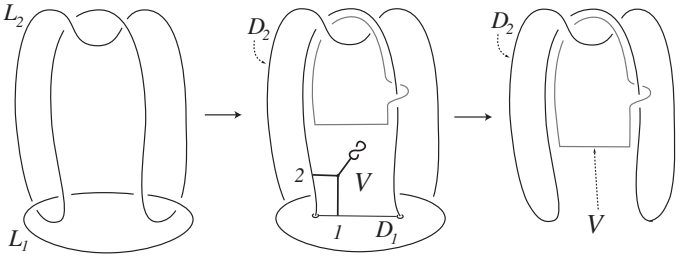

FIGURE 3. *The Whitehead link L of Figure 2 also bounds an order 2 twisted Whitney tower $\mathcal{V}$ consisting of embedded disks D_1 and D_2 bounded by L_1 and L_2, together with an embedded twisted Whitney disk V pairing the intersections $D_1 \cap D_2$. Moving into B^4 from left to right, the left picture again shows $L = L_1 \cup L_2 \subset S^3$, and disks D_1 and D_2 are described just as in Figure 2, with D_1 contained in the left and middle pictures, while circle slices of D_2 are visible in all the three pictures with the rest of D_2 (not shown) extending further into B^4 as an embedded disk bounding the indicated unlink component in the right-hand picture. The embedded twisted Whitney disk V containing its associated twisted tree is partly visible in the middle picture, which shows an opaque annular region of V that contains the boundary of V. The rest of V extends further into B^4 and is visible as the indicated component of the unlink in the right picture together with an embedded disk (not shown) bounding this component which is disjoint from the part of D_2 that bounds the other unlink component. That V is twisted will be shown in Section 6.*

In [**10**], we studied the *twisted Whitney tower filtration*:

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

where a link L lies in $\mathbb{W}_n^\infty$ if and only if L bounds a twisted Whitney tower $\mathcal{W}$ in the 4-ball which is of *order* n, meaning very roughly that $\mathcal{W}$ consists of n layers of Whitney disks on top of the immersed disks bounding the components of L; for details see Section 2. The 'twist' symbol ∞ in our notation stands for the fact that some of the Whitney disks in a twisted Whitney tower are allowed to be *twisted* (that is, allowed to have non-zero relative normal Euler number), as opposed to being *framed*.

THEOREM 1. *There is an inclusion $\mathbb{W}_n^\infty \subseteq \mathbb{M}_n$ of filtrations. Moreover, all length $(n+2)$ Milnor invariants of $L \in \mathbb{W}_n^\infty$ are defined and can be computed from the intersection invariant $\tau_n^\infty(\mathcal{W})$ of any order n twisted Whitney tower $\mathcal{W}$ bounded by L.*

The second statement will be made precise in Theorem 6, which describes exactly how Milnor invariants are measured by the higher-order intersections and Whitney disk twistings that determine the invariant τ_n^∞. The theorem also works for an order 0 twisted Whitney tower $\mathcal{W}$, which by definition is just a union of immersed disks (which are oriented consistently with the orientation of L): $\tau_0^\infty(\mathcal{W})$ counts the transverse intersections between pairs of those disks, well known to equal the linking numbers $\mu_L(i,j)$ of the link L, and also detects the induced framings on the link components, considered as 'self-linking' numbers $\mu_L(i,i)$ (see the proof of Theorem 6 in Section 4).

Any L in $\mathbb{M}_1$ bounds an order 1 twisted Whitney tower consisting of immersed disks D_i together with Whitney disks pairing all intersections among D_i. Matsumoto [**35**] showed that $\mu_L(i,j,k)$ can be computed from the interior intersections between the Whitney disks and the D_i (see Paragraph 4.2.1 and Figure 13). Figure 1 shows the easiest case where one can explicitly see why a Whitney move cannot be used to find three disjointly embedded disks in B^4 whose boundaries form the Borromean rings.

Our results generalize the correspondence between Milnor invariants and higher-order intersection invariants of Whitney towers to all lengths and all orders. We find some surprising subtleties related to Whitney disk twistings, with the first occurrence in order $n = 2$, as illustrated in Subsection 1.4 (see Figure 3).

To express more precisely the difference between the two filtrations $\mathbb{M}_n$ and $\mathbb{W}^\infty_n$, we work with the associated graded *groups*. More precisely, in Subsection 1.1 we present a *universal* order n Milnor invariant $\mu_n : \mathbb{M}_n \to \mathsf{D}_n$ with values in a free abelian group of known rank. It carries exactly the information of all Milnor invariants of length $(n + 2)$ and can be expressed in the language of trivalent trees that are closely related to the intersection invariant τ^∞_n. Moreover, it has the additivity properties

$$\mu_n(L\#_b L') = \mu_n(L) + \mu_n(L') \quad \text{and} \quad \mu_n(-L) = -\mu_n(L),$$

where $L\#_b L'$ is any choice of band connected sum of m-component links $L, L' \in \mathbb{M}_n$ (that are separated by an embedded 2-sphere) and $-L$ denotes the link L mirror reflected and with orientations flipped. This additivity follows most easily from the translation of Milnor invariants into Massey products for the link complement, shown independently by Turaev [**52**] and Porter [**41**].

Since $\mathbb{M}_{n+1}$ consists exactly of those links $L \in \mathbb{M}_n$ with $\mu_n(L) = 0$, we think of group D_n as being in some sense the 'quotient group' of $\mathbb{M}_n$ by $\mathbb{M}_{n+1}$. The same band connected sum operation on links makes $\mathbb{W}^\infty_n$ into a finitely generated abelian group W^∞_n after taking the quotient by the equivalence relation of *order* $n + 1$ *twisted Whitney tower concordance* [**10**], and the following theorem gives 'three quarters' of our main classification result.

THEOREM 2. *The universal Milnor invariants induce surjections* $\mu_n \colon \mathsf{W}^\infty_n \twoheadrightarrow \mathsf{D}_n$ *for all* n, *and isomorphisms for all* $n \equiv 0, 1, 3 \bmod 4$.

This result follows from Theorem 1 via [**10**, **11**], as will be explained in Subsection 1.4 (see Corollary 7 and Theorem 8). For the remaining 'quarter' of orders, we need certain *higher-order Arf invariants* Arf_k, as described in [**9**, **10**] and sketched below in Subsection 1.5. Arf_k are link concordance invariants which represent obstructions to *framing* a twisted Whitney tower bounded by a link, with Arf_1 corresponding to the classical Arf invariants of the link components. Although the precise values of Arf_k are not known for $k > 1$ (they form a quotient of a known finite 2-torsion group, see Definition 11), we show in this paper that these invariants measure the difference between two natural notions of 'nilpotent approximation' of slice disks for a link: *k-slice* and *geometrically k-slice* (Theorem 16). The construction of boundary links realizing the range of Arf_k (Lemma 13) yields two new geometric characterizations of links with vanishing length at most $2k$ Milnor invariants, as described in Theorems 17 and 18.

The rest of this introduction develops enough material to give more refined statements of our results.

1.1. *Order n Milnor invariants*

For a group Γ, denote by Γ_n the *lower central series* of commutator subgroups of Γ, defined inductively by $\Gamma_1 := \Gamma$ and $\Gamma_{n+1} := [\Gamma, \Gamma_n]$. If $L \subset S^3$ is an m-component link such that all of its longitudes lie in the $(n + 1)$th term of the lower central series of the link group $\pi_1(S^3 \setminus L)_{n+1}$, then a choice of meridians induces an isomorphism

$$\frac{\pi_1(S^3 \setminus L)_{n+1}}{\pi_1(S^3 \setminus L)_{n+2}} \cong \frac{F_{n+1}}{F_{n+2}},$$

where $F = F(m)$ is the free group on $\{x_1, x_2, \ldots, x_m\}$.

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

In this setting, denote by $\mu_n^i(L)$ the image of the ith longitude in L_{n+1} under the above isomorphisms and define the *order n Milnor invariant* $\mu_n(L)$ by

$$\mu_n(L) := \sum_{i=1}^{m} X_i \otimes \mu_n^i(L) \in \mathsf{L}_1 \otimes \mathsf{L}_{n+1}.$$

Then $\mu_n(L)$ is the first non-vanishing *universal* Milnor invariant, in the sense that it determines *all* Milnor invariants of *length* $n+2$ (with repeating indices allowed) [**37**, **38**]. The original $\bar{\mu}$-invariants are computed from the longitudes via the Magnus expansion as integers modulo indeterminacies coming from invariants of shorter length. Since in this paper we are only concerned with first non-vanishing μ-invariants, we do not use the 'bar' notation $\bar{\mu}$.

It turns out that $\mu_n(L)$ actually lies in the kernel $\mathsf{D}_n = \mathsf{D}_n(m)$ of the bracket map $\mathsf{L}_1 \otimes \mathsf{L}_{n+1} \to \mathsf{L}_{n+2}$ (for example, by 'cyclic symmetry' [**19**]). We observe that L_n and D_n are free abelian groups of known ranks: The rank $r_n = r_n(m)$ of $\mathsf{L}_n(m)$ is given by $r_n = (1/n)\sum_{d|n} M(d) m^{n/d}$, with M denoting the Möbius function [**34**, Theorem 5.11]; and the rank of $\mathsf{D}_n(m)$ is equal to $mr_{n+1} - r_{n+2}$, first identified by Orr as the number of independent (integer) μ-invariants of length $n+2$ in [**40**].

Milnor's *$\bar{\mu}$-invariants* have inspired a significant amount of research over the past 50-plus years. Work of Stallings [**49**] implied that Milnor invariants are concordance invariants [**2**]. Realization results of Cochran [**5**, **6**] and Orr [**40**] provided geometric interpretations of $\bar{\mu}$-invariants [**5**, **27**] and supported the development of more 'universal' approaches, including Habbeger and Lin's classification of link homotopy [**21**] via an Artin representation characterization of Milnor invariants for string links (see also [**22**]) as well as a growing number of interpretations related to quantum invariants (for example, [**1**, **7**, **23**–**25**, **36**]). There are even connections with algebraic number theory [**39**] and molecular biology [**20**]. See, for example, [**26**, Chapter 10] for more regarding Milnor invariants.

1.2. *Intersection invariants for (twisted) Whitney towers*

By Conant, Schneiderman and Teichner [**10**] and Schneiderman and Teichner [**47**], an order n (twisted) Whitney tower $\mathcal{W}$ built on properly immersed disks in the 4-ball has an *intersection invariant* $\tau_n(\mathcal{W})$ (respectively, $\tau_n^{\infty}(\mathcal{W})$), which is defined by associating unitrivalent trees to the unpaired higher-order intersection points (and twisted Whitney disks) in $\mathcal{W}$ (for example, Figures 2 and 3). For the convenience of the reader, we briefly describe next the target groups of these invariants. Relevant details on (twisted) Whitney towers and their intersection invariants are presented in Section 2.

DEFINITION 3. In this paper, a *tree* will always refer to a finite oriented unitrivalent 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 $\{1, 2, 3, \ldots, m\}$ 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 following antisymmetry (AS) and Jacobi (IHX) relations

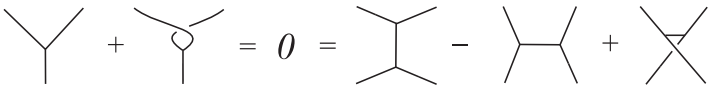

Since the AS and IHX relations are homogeneous with respect to order, $\mathcal{T}$ inherits a grading $\mathcal{T} = \bigoplus_n \mathcal{T}_n$, where $\mathcal{T}_n = \mathcal{T}_n(m)$ is the free abelian group on order n trees, modulo AS and IHX relations.

In the Whitney tower obstruction theory of [47], the order n intersection invariant $\tau_n(\mathcal{W}) \in \mathcal{T}_n$ assigned to each order n (framed) Whitney tower $\mathcal{W}$ is defined by summing the trees associated to unpaired intersections in $\mathcal{W}$ (see Figure 2 for an example). The tree orientations are induced by Whitney disk orientations via a convention that corresponds to the AS relations (Subsection 2.5), and the IHX relations can be realized geometrically by controlled maneuvers on Whitney towers as described in [8, 44]. It follows from the obstruction theory that a link bounds an order n Whitney tower $\mathcal{W}$ with $\tau_n(\mathcal{W}) = 0$ if and only if it bounds an order $n+1$ Whitney tower [47, Theorem 2].

For twisted Whitney towers of order n, the intersection invariant τ_n^{∞} introduced in [10] also assigns certain *twisted trees* (∞-trees) to Whitney disks which are not framed (see Subsection 2.3), and takes values in the following graded groups.

DEFINITION 4 [10]. In odd orders, the group $\mathcal{T}^{\infty}_{2k-1}$ is the quotient of $\mathcal{T}_{2k-1}$ by the *boundary-twist relations*:

$$i \multimap \overset{J}{<} = 0,$$

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

A *rooted* tree has one unlabeled univalent vertex, called the *root*. For any rooted tree J, we define the corresponding *∞-tree*, denoted by J^{∞}, by labeling the root univalent vertex with the twist-symbol '∞':

$$J^{\infty} := \infty \multimap J.$$

In even orders, group $\mathcal{T}^{\infty}_{2k}$ is the free abelian group on trees of order $2k$ and ∞-trees of order k, modulo the following four types of relations:

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

Here the *inner product* $\langle T_1, T_2\rangle$ of two order k rooted trees T_1 and T_2 is defined by gluing the roots together to get an unrooted tree of order $2k$. The AS and IHX relations are as pictured above, but they only apply to ordinary trees (not to ∞-trees).

These relations have the following geometric origins: The *symmetry relation* corresponds to the fact that the relative Euler number of a Whitney disk is independent of its orientation, with the minus sign denoting that the cyclic edge-orderings at the trivalent vertices of $-J$ differ from those of J at an odd number of vertices. As explained in [10], the *twisted IHX relation* corresponds to the effect of performing a Whitney move in the presence of a twisted Whitney disk, and the *interior-twist relation* corresponds to the fact that creating a ± 1 self-intersection in a Whitney disk changes its twisting by ∓ 2.

The main reason why this is a good definition comes from the following result, which is [10, Theorem 2.10]: A link $L \subset S^3$ bounds an order n twisted Whitney tower $\mathcal{W} \subset B^4$ with $\tau_n^{\infty}(\mathcal{W}) = 0 \in \mathcal{T}_n^{\infty}$ if and only if L bounds an order $n+1$ twisted Whitney tower in B^4.

1.3. *The summation maps η_n*

The connection between $\tau_n^{\infty}(\mathcal{W})$ and $\mu_n(L)$ is via a homomorphism $\eta_n : \mathcal{T}_n^{\infty} \to \mathsf{D}_n$, which is best explained when we regard rooted trees of order n as elements in L_{n+1} in the usual way: For v a univalent vertex of an order n tree t as in Definition 3, denote by $B_v(t) \in \mathsf{L}_{n+1}$ the Lie

bracket of generators $X_1, X_2, \dots, X_m$ determined by the formal bracketing of indices which is gotten by considering v to be a root of t.

DEFINITION 5. Denoting the label of a univalent vertex v by $\ell(v) \in \{1, 2, \dots, m\}$, the map $\eta_n : \mathcal{T}_n^{\infty} \to \mathsf{L}_1 \otimes \mathsf{L}_{n+1}$ is defined on generators by

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

The first sum is over all univalent vertices v of t, and the second expression lies in $\mathsf{L}_1 \otimes \mathsf{L}_{n+1}$ because the coefficients of $\eta_n(\langle J, J \rangle)$ are even. Here J is a rooted tree of order k for $n = 2k$.

For example,

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

And similarly,

$$\begin{aligned}\eta_2\big(\infty \multimap \tfrac{2}{1}\big) &= \frac{1}{2}\,\eta_2\big(\tfrac{1}{2} \succ\!\prec \tfrac{2}{1}\big) \\ &= X_1 \otimes 2 \succ\!\prec \tfrac{2}{1} + X_2 \otimes 1 \succ\!\prec \tfrac{2}{1} \\ &= X_1 \otimes [X_2, [X_1, X_2]] + X_2 \otimes [[X_1, X_2], X_1].\end{aligned}$$

In Subsection 4.3, we check that η_n is well defined and maps $\mathcal{T}_n^{\infty}$ onto D_n.

We can now make Theorem 1 precise as follows.

THEOREM 6. *If L bounds a twisted Whitney tower $\mathcal{W}$ of order n, then the order k Milnor invariants $\mu_k(L)$ vanish for $k < n$ and*

$$\mu_n(L) = \eta_n \circ \tau_n^{\infty}(\mathcal{W}) \in \mathsf{D}_n.$$

The proof of Theorem 6 given in Section 4 uses a geometric interpretation of the maps η_n which shows precisely how higher-order intersections and Whitney disk twistings correspond to sums of iterated commutators. Closely related maps over the rationals appear in Habegger and Masbaum's work on the Kontsevich invariant showing that the Milnor invariants are the universal finite type (rational) concordance invariants [**23**]. Levine's work on homology cylinders [**31**–**33**] led him to study analogous maps over the integers, and our resolution in [**11**] of his main conjecture accomplished an important step in the classification of both the twisted and framed Whitney tower filtrations on link concordance, as discussed in the next two subsections. It should also be noted that the relationship between Milnor invariants and trees originally goes back to Cochran's method of constructing links realizing given (integer) Milnor invariants [**5**, **6**].

1.4. *Computing the graded groups associated to the filtrations*

As a prelude to the description of our results on the higher-order Arf invariants, we briefly recall from [**10**] the computation of the groups W_n^{∞} and W_n associated to the twisted and framed Whitney tower filtrations on link concordance. Here the groups W_n in the framed setting are defined analogously to W_n^{∞}: The equivalence relation of (framed) *Whitney tower concordance of order $n+1$* on the set $\mathbb{W}_n$ of m-component links that bound (framed) order n Whitney towers in the 4-ball defines W_n as the associated graded quotient, which also turns out to be a finitely generated abelian group under the band connected sum operation for all n.

In [10], we constructed framed and twisted *realization* epimorphisms

$$R_n : \mathcal{T}_n \twoheadrightarrow \mathsf{W}_n \quad \text{and} \quad R_n^{\infty} : \mathcal{T}_n^{\infty} \twoheadrightarrow \mathsf{W}_n^{\infty},$$

which send $g \in \mathcal{T}_n^{(\infty)}$ to the equivalence class of links bounding an order n (twisted) Whitney tower $\mathcal{W}$ with $\tau_n^{(\infty)}(\mathcal{W}) = g$. These maps are defined similarly to Cochran's construction for realizing Milnor invariants [5, Section 7; 6, Theorem 3.3] by 'Bing-doubling along trees' and taking internal band sums if indices repeat: The Hopf link realizes the order 0 tree $1 \text{——} 2$ corresponding to a transverse intersection between disks in B^4 bounded by the components. To realize higher-order generators, iterated (untwisted) Bing-doublings are performed according to the branching of the tree, until we obtain the correct tree but with non-repeating indices labeling the univalent vertices. For example, a single Bing-doubling on the Hopf link yields the Borromean rings realizing $1 \text{—}\!\prec \frac{3}{2}$. Finally, we take internal band sums according to which indices repeat. For example, one internal band sum may take the Borromean rings to the Whitehead link defining $R_1(1 \text{—}\!\prec \frac{2}{2})$ (see Figure 2).

For ∞-trees, the starting point is the 1-framed unknot as $R_0^{\infty}(\infty \text{—} 1)$. The first Bing-doubling has to be a twisted one, giving a Whitehead link as $R_2^{\infty}(\infty \text{—}\!\prec \frac{1}{2})$. Note that this means that the Whitehead link bounds two different Whitney towers, one framed of order 1 and the other twisted of order 2 (Figure 3). This is the easiest example illustrating how the Milnor invariants interact differently with the framed and twisted Whitney tower filtrations: The Whitehead link L bounds a twisted Whitney tower $\mathcal{V}$ of order 2, but not one of order 3, as detected by $\tau_2^{\infty}(\mathcal{V}) = \infty \text{—}\!\prec \frac{1}{2} \neq 0 \in \mathcal{T}_2^{\infty}$, corresponding to the non-triviality of $\mu_2(L)$. However, even though the longitudes of L are length-3 commutators (so that $\mu_2(L)$ is defined), L does not bound an order 2 *framed* Whitney tower; as detected by $\tau_1(\mathcal{W}) = 1 \text{—}\!\prec \frac{2}{2} \neq 0 \in \mathcal{T}_1$, corresponding to the non-trivial *Sato–Levine invariant* [43] which is the projection of $\mu_2(L)$. This phenomenon occurs in all odd orders of the framed Whitney tower filtration, as described by the *higher-order Sato–Levine invariants* defined in [10].

The maps R_n^{∞} bound the size of the abelian groups W_n^{∞} from above, and the following corollary of Theorem 6 shows that Milnor invariants give a lower bound.

COROLLARY 7. *There is a commutative diagram of epimorphisms*

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

As a consequence of our proof [11] of the combinatorial conjecture of Levine formulated in [32], we have the following partial classification of W_n^{∞}.

THEOREM 8 [10]. *The maps $\eta_n : \mathcal{T}_n^{\infty} \to \mathsf{D}_n$ are isomorphisms for $n \equiv 0, 1, 3 \bmod 4$. As a consequence, both the Milnor invariants $\mu_n \colon \mathsf{W}_n^{\infty} \to \mathsf{D}_n$ and the twisted realization maps $R_n^{\infty} : \mathcal{T}_n^{\infty} \to \mathsf{W}_n^{\infty}$ are isomorphisms for these orders.*

The remaining cases to complete the classification are more complicated, as can already be seen for $n = 2$: In the case $m = 1$ of knots, Lemma 10 shows that the Arf invariant induces an isomorphism $\mathsf{W}_2^{\infty}(1) \cong \mathbb{Z}_2$, whereas all Milnor invariants vanish for knots.

Unlike for $n \equiv 0, 1, 3 \bmod 4$, where $\operatorname{Ker}(\eta_n) = 0$, there are some obvious elements in $\operatorname{Ker}(\eta_{4k-2})$, namely those of the form $\infty \text{—}\!\prec \frac{J}{J}$ for an order $k-1$ rooted tree J: These are

2-torsion by the interior-twist and IHX relations in $\mathcal{T}^{\omega}_{4k-2}$ and hence must map to zero in the (torsion-free) group D_{4k-2}. In [10], we also deduce the following result from the affirmation of Levine's conjecture.

PROPOSITION 9 [10]. *The map sending* $1 \otimes J$ *to* $\omega \multimap^{J} \in \mathcal{T}^{\omega}_{4k-2}$ *for rooted trees* J *of order* $k-1$ *defines an isomorphism*:

$$\mathbb{Z}_2 \otimes \mathsf{L}_k \cong \operatorname{Ker}(\eta_{4k-2}).$$

Here the identification of rooted order $k-1$ trees with degree k Lie brackets is as in Subsection 1.3 (see the examples following Definition 5). It follows that $\mathbb{Z}_2 \otimes \mathsf{L}_k$ is also an upper bound on the kernels of the epimorphisms $R^{\omega}_{4k-2} : \mathcal{T}^{\omega}_{4k-2} \twoheadrightarrow \mathsf{W}^{\omega}_{4k-2}$ and $\mu_{4k-2} : \mathsf{W}^{\omega}_{4k-2} \twoheadrightarrow \mathsf{D}_{4k-2}$, and the calculation of $\mathsf{W}^{\omega}_{4k-2}$ is completed by invariants defined on the kernel of μ_{4k-2} which are the above-mentioned higher-order Arf invariants, as we describe next.

1.5. *Higher-order Arf invariants*

Let us first discuss the situation for order $n = 2$. Observe that $\omega \multimap^{1}_{1}$ is not zero in $\mathcal{T}^{\omega}_{2}(1)$ but that $\overset{1}{1} \rightarrowtail\!\!\!\!\multimap^{1}_{1} = 0$ by the IHX relation; so $\mathcal{T}^{\omega}_{2}(1)$ is generated by $\omega \multimap^{1}_{1}$, which is 2-torsion, and $\tau^{\omega}_{2}(\mathcal{W})$ counts (modulo 2) the framing obstructions on the Whitney disks in an order 2 twisted Whitney tower $\mathcal{W}$. This is explained in Section 5, which gives a proof of the following result.

LEMMA 10. *Any knot* K *bounds a twisted Whitney tower* $\mathcal{W}$ *of order* 2 *and the classical Arf invariant of* K *can be identified with the intersection invariant*

$$\tau^{\omega}_{2}(\mathcal{W}) \in \mathcal{T}^{\omega}_{2}(1) \cong \mathbb{Z}_2.$$

More generally, the classical Arf invariants of the components of an m*-component link give an isomorphism*

$$\operatorname{Arf} : \operatorname{Ker}(\mu_2 : \mathsf{W}^{\omega}_{2} \twoheadrightarrow \mathsf{D}_2) \xrightarrow{\cong} (\mathbb{Z}_2 \otimes \mathsf{L}_1) \cong (\mathbb{Z}_2)^m.$$

This lemma verifies our conjecture $\mathsf{W}^{\omega}_{n} \cong \mathcal{T}^{\omega}_{n}$ from Conjecture 12 in the case $n = 2$, with $\operatorname{Ker}(\eta_2) \cong \operatorname{Ker}(\mu_2) \cong (\mathbb{Z}_2)^m$.

Following and expanding upon [10], we will now describe a similarly satisfying picture for all orders of the form $n = 4k-2$ that takes both the Milnor and Arf invariants into account.

Let $\mathsf{K}^{\omega}_{4k-2}$ denote the kernel of μ_{4k-2}. It follows from Corollary 7 and Proposition 9 that mapping $1 \otimes J$ to $R^{\omega}_{4k-2}(\omega \multimap^{J})$ induces a surjection $\alpha^{\omega}_{k} : \mathbb{Z}_2 \otimes \mathsf{L}_k \twoheadrightarrow \mathsf{K}^{\omega}_{4k-2}$, for all $k \geqslant 1$. Denote by $\overline{\alpha^{\omega}_{k}}$ the induced isomorphism on $(\mathbb{Z}_2 \otimes \mathsf{L}_k)/\operatorname{Ker}\alpha^{\omega}_{k}$.

DEFINITION 11 [10]. The *higher-order Arf invariants* are defined by

$$\operatorname{Arf}_k := (\overline{\alpha^{\omega}_{k}})^{-1} : \mathsf{K}^{\omega}_{4k-2} \longrightarrow (\mathbb{Z}_2 \otimes \mathsf{L}_k)/\operatorname{Ker}\alpha^{\omega}_{k}.$$

From Theorem 8, Proposition 9 and Definition 11, we see that the groups W^{ω}_{n} are computed by the Milnor and higher-order Arf invariants.

We conjectured in [9, 10] that α^{ω}_{k} is an isomorphism, which would mean that Arf_k are very interesting new concordance invariants.

CONJECTURE 12 [9, 10]. The map $\operatorname{Arf}_k : \mathsf{K}^{\omega}_{4k-2} \to \mathbb{Z}_2 \otimes \mathsf{L}_k$ is an isomorphism for all k.

Conjecture 12 would imply that

$$\mathsf{W}^{\infty}_{4k-2} \cong \mathcal{T}^{\infty}_{4k-2} \cong (\mathbb{Z}_2 \otimes \mathsf{L}_k) \oplus \mathsf{D}_{4k-2},$$

where the second isomorphism (is non-canonical and) already follows from Corollary 7 and Proposition 9 [**10**, Corollary 1.12, Proposition 1.14]. The statement of Conjecture 12 is true for $k = 1$, by Lemma 10, with $\mathrm{Arf}_1 = \mathrm{Arf}$ the classical Arf invariant. It remains an open problem whether Arf_k is non-trivial for any $k > 1$.

We have the following specialization of the Bing-doubling construction discussed above Corollary 7, which realizes symmetric ∞-trees of the form $\infty \!\!-\!\!\prec^{J}_{J}$. It starts with the fact that any knot with non-trivial Arf invariant represents $R^{\infty}_2(\infty \!\!-\!\!\prec^{1}_{1})$ by Lemma 10, and then proceeds by applying untwisted Bing-doublings and internal band sums. This has the effect of symmetrically extending both 1-labeled branches of the original tree into a higher-order twisted Whitney tower, so that the resulting link can be constructed to realize any $\infty \!\!-\!\!\prec^{J}_{J}$ (see Subsection 6.1). This idea will be used to derive the geometric interpretations of Milnor invariants given in Theorems 17 and 18.

LEMMA 13. *Let J be any rooted tree of order $k-1$. By performing iterated untwisted Bing-doublings and interior band sums on the figure-eight knot K, a boundary link K^J can be constructed which bounds a twisted Whitney tower $\mathcal{W}$ of order $4k-2$ such that*

$$\tau^{\infty}_{4k-2}(\mathcal{W}) = \infty \!\!-\!\!\prec^{J}_{J} .$$

The proof of Lemma 13 given in Subsection 6.1 can be easily modified to show that this result holds for any knot K with non-trivial classical Arf invariant. It is thus already interesting to ask whether our proposed higher-order Arf invariants Arf_k can be defined on the cobordism group of boundary links. The links K^J of Lemma 13 are known not to be slice by work of Cha [**4**], providing evidence supporting our conjecture that Arf_k is indeed a non-trivial link concordance invariant which represents an obstruction to bounding an order $4k-1$ twisted Whitney tower. The following result emphasizes the importance of the first open case $k = 2$.

PROPOSITION 14. *If Arf_2 is trivial, then Arf_k is trivial for all $k \geqslant 2$.*

As explained in Subsection 6.2, which contains a proof of Proposition 14, the statement that Arf_2 is trivial is equivalent to the existence of an order 7 twisted Whitney tower $\mathcal{W}$ bounded by the Bing-double of a figure-eight knot.

1.6. *Geometrically k-slice links*

We conclude this introduction with some new geometric characterizations of Milnor invariants and the higher-order Arf invariants. See Section 7 for proofs of these results.

Recall (for example, from [**51**]) that the 2-complexes known as *gropes* are 'geometric embodiments' of iterated commutators in the sense that a loop in a topological space represents a k-fold commutator in the fundamental group if and only if it extends to a continuous map of a grope of class k (see Section 3 and for example, [**8**, **12**, **13**, **18**, **19**, **29**, **30**, **45**, **50**]). Since Milnor invariants measure how deeply the link longitudes extend into the lower central series of the link group, Milnor invariants obstruct bounding *immersed* gropes in S^3 essentially by definition. On the other hand, extracting information on bounding *embedded* gropes in the 4-ball from the vanishing of Milnor invariants is much more difficult. Embedded framed gropes have usefully served as 'approximations' to embedded disks in many topological settings (see, for example, [**51**]).

Perhaps the most notable geometric 'if and only if' characterization of Milnor invariants to date is the *k-slice Theorem*, due to Igusa and Orr: Expressed in the language of gropes, a link $L \subset S^3$ is said to be *k-slice* if the link components L_i bound disjointly embedded (oriented) surfaces $\Sigma_i \subset B^4$ such that a symplectic basis of curves on each Σ_i bound class k gropes immersed in the complement of $\Sigma := \bigcup_i \Sigma_i$. Via a very careful analysis of the third homology of the nilpotent quotients F/F_k of the (rank m) free group F, Igusa and Orr [**27**] proved the following result.

THEOREM 15 [**27**]. *A link L is k-slice if and only if $\mu_n(L) = 0$ for all $n \leqslant 2k-2$ (equivalently, all Milnor invariants of length at most $2k$ vanish).*

The k-slice condition says that the link components bound certain immersed gropes in B^4 whose embedded bottom stage surfaces are 'algebraic approximations' of slice disks modulo the kth term of the lower central series of the link group.

This leads to the very natural notion of *geometrically k-slice* links: These are links for which there is a symplectic basis of curves on the embedded bounding surfaces $\Sigma \subset B^4$ that bound *disjointly embedded framed* gropes of class k in $B^4 \setminus \Sigma$. In Section 7, we describe how the techniques of [**45**] together with the classification of the twisted Whitney tower filtration in [**10**] can be used to give the following result, which shows that the higher-order Arf invariants Arf_k detect the difference between k-sliceness and geometric k-sliceness.

THEOREM 16. *A link L is geometrically k-slice if and only if $\mu_n(L) = 0$ for all $n \leqslant 2k-2$ and $\mathrm{Arf}_n(L) = 0$ for all $n \leqslant k$.*

Combining Theorem 16 together with Corollary 7, Proposition 9 and Lemma 13, we immediately get the following theorem.

THEOREM 17. *A link L has $\mu_n(L) = 0$ for all $n \leqslant 2k-2$ if and only if L is geometrically k-slice after a finite number of band sums with boundary links.* □

It turns out that the operation of taking band sums with boundary links is equivalent to a certain approximation of being geometrically k-slice, as described by the following theorem. The basic observation here is that any curve on a surface in S^3 bounds an immersed disk in B^4, leading to the surfaces of type Σ_i'' below, associated to the boundary links in Theorem 17. We note that the 'only if' part of the following theorem uses a mild generalization of Theorem 6, described in Proposition 34.

THEOREM 18. *A link $L = \bigcup_i L_i$ has $\mu_n(L) = 0$ for all $n \leqslant 2k-2$ if and only if the link components L_i 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 $\Sigma := \bigcup_i \Sigma_i$ 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 18 is a considerable strengthening of the above Igusa–Orr k-slice theorem: 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

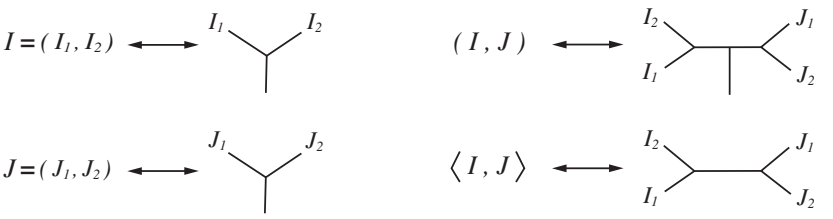

FIGURE 4. *The rooted product* (I, J) *and inner product* $\langle I, J\rangle$ *of* $I = (I_1, I_2)$ *and* $J = (J_1, J_2)$. *All trivalent orientations correspond to a clockwise orientation of the plane.*

class k found by Igusa and Orr can be cleaned up to immersed *disks* (these are immersed gropes of arbitrarily high class) or *disjointly embedded* gropes of class k. As explained in Section 7, the higher-order Arf invariants are exactly the obstructions to eliminating the need for the Σ_i'' and these immersed disks.

2. *Whitney towers*

This section recalls the relevant theory of (twisted) Whitney towers as developed in [**8**, **10**, **45**, **47**]. 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 [**18**]. In fact, it can be shown that the filtrations $\mathbb{W}_n$ and $\mathbb{W}_n^\infty$ 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 [**10**, Remark 2.1]).

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 as in Definition 3, with each univalent vertex labeled by an element from the index set, except for the *root* univalent vertex which is left unlabeled.

DEFINITION 19. 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 4 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 4 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 [**45**, **47**]). In [**45**, **47**], the inner product is written as a dot-product, and the rooted product is denoted by $*$.

2.2. *Whitney disks and higher-order intersections*

DEFINITION 20. 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

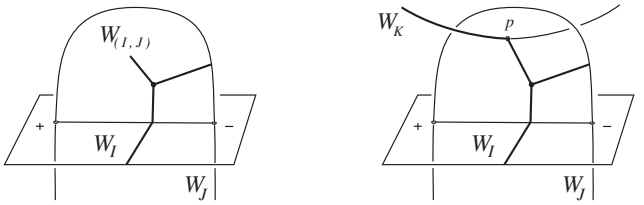

FIGURE 5. *On the left, (part of) the rooted tree (I, J) associated to a Whitney disk $W_{(I,J)}$. On the right, (part of) the tree $t_p = \langle (I, J), K\rangle$ associated to an intersection $p \in W_{(I,J)} \cap W_K$. Note that p corresponds to where the roots of (I, J) and K are identified to a (non-vertex) point in $\langle (I, J), K\rangle$.*

boundary is embedded in ∂M and the interior is generically immersed in $M \smallsetminus \partial M$. The A_i are also required to be *framed* as discussed in Subsection 2.3.

To each *order* 0 *surface* A_i is associated the order 0 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 0 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 Subsection 2.5.)

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 5); with the understanding that if, say, I is just a singleton i, then W_I denotes the order 0 surface A_i.

To any transverse intersection $p \in W_{(I,J)} \cap W_K$ between $W_{(I,J)}$ and any W_K is associated the tree $t_p := \langle (I, J), K\rangle$ (see right-hand side of Figure 5).

DEFINITION 21. 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 22. 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 0 surfaces must also be *framed*, as discussed next.

2.3. *Twisted Whitney disks and framings*

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 that can be described in the following way: The Whitney disk boundary circle ∂W is the union of two arcs, each lying in a local sheet of the surfaces paired by W. The Whitney section is given by pushing ∂W tangentially along one sheet and normally off of the other sheet (while avoiding the tangential direction of W). See Figure 6, and for more details, for example, [**48**, 1.7]. Pulling

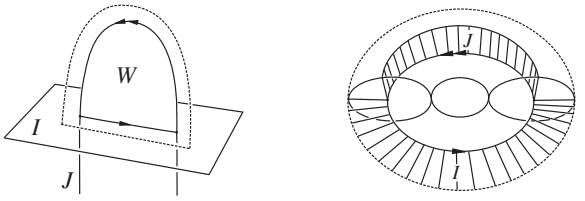

FIGURE 6. *In a three-dimensional slice of* 4*-space* (*left*), *the Whitney section over the boundary of a framed Whitney disk* W *is indicated by the dotted loop, where* ∂W *has been pushed tangentially along the* I*-sheet and normally off of the* J*-sheet. On the right is shown an embedding into* 3*-space of the normal disk-bundle to* W *over* ∂W (*a solid torus, shown mostly transparent*), *with the dotted loop again indicating the Whitney section and the* '*thatched*' *lines indicating parts of the surface sheets. The* I*-labeled thatches indicate the tangential push of* ∂W *along the* I*-sheet, and the* J*-labeled thatches indicate* '*one side*' *of the* J*-sheet along* ∂W.

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.

Note that for order 0 surfaces 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 M is oriented, this is equivalent to an orientation and a non-vanishing normal vector field on ∂A_i (respectively, A_i). The twisting $\omega(A_i) \in \mathbb{Z}$ of an order 0 surface is also defined when a framing of ∂A_i is given, and the order 0 surface A_i is said to be *framed* when $\omega(A_i) = 0$.

2.4. *Twisted Whitney towers*

In the definition of an order n Whitney tower given just above (following [8, 44, 45, 47]), all Whitney disks and order 0 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 23 [10]. A *twisted Whitney tower of order* 0 is a Whitney tower of order 0 without any framing requirement (a collection of properly immersed surfaces in a 4-manifold).

For $k > 0$, a *twisted Whitney tower of order* $2k-1$ is just a (framed) Whitney tower of order $2k-1$ as in Definition 22.

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

REMARK 24. 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 23.

REMARK 25. The convention of allowing only order at least k twisted Whitney disks in order $2k$ twisted Whitney towers will be explained in Section 4, where it will be seen that

twisted Whitney disks contribute to the link longitudes just as described by the definition of the η-map on ∞-trees.

In any event, an order $2k$ twisted Whitney tower can always be modified so that all its Whitney disks of order greater than k are framed, so the twisted Whitney disks of order equal to k are the ones relevant to the obstruction theory [**10**, Section 4.1].

2.5. *Whitney tower orientations*

Orientations on order 0 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 5. (In fact, if a tree t has more than one trivalent vertex corresponding 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 5.)

This 'negative corner' convention, which differs from the positive corner convention in the earlier papers [**8**, **47**] but agrees with all more recent papers on Whitney towers, will turn out to be compatible with commutator conventions for use in Section 4.

With these conventions, different choices of orientations on Whitney disks in $\mathcal{W}$ correspond to AS relations (as explained in [**47**, Section 3.4]).

2.6. *Links bounding (twisted) Whitney towers*

Throughout this paper, the statement that a link $L \subset S^3$ bounds an order n (twisted) Whitney tower $\mathcal{W} \subset B^4$ means that the components of L bound properly immersed disks which are the order 0 surfaces of $\mathcal{W}$ as in Definition 22 (Definition 23), with all conventions as described above.

2.7. *Intersection invariants for twisted Whitney towers*

The intersection invariants for twisted Whitney towers take values in the groups $\mathcal{T}_n^\infty$ defined in the introduction (Definition 4).

Recall from Definition 23 (and Remark 25) 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 26 [**10**]. The *order n 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 Subsection 2.2 our notational convention that W_j denotes A_j; in this case, $\omega(A_j) \in \mathbb{Z}$ is the relative Euler number of the normal bundle of A_j with respect to the given framing of ∂A_j as in Subsection 2.3.

By splitting the twisted Whitney disks, as explained in Subsection 2.8, for $n > 0$ we may actually assume that all non-zero $\omega(W_J) \in \{\pm 1\}$, just like the signs ϵ_p.

The vanishing of τ_n^∞ is sufficient for the existence of a twisted Whitney tower of order $(n+1)$.

THEOREM 27 [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.*

2.8. *Split twisted Whitney towers*

A twisted Whitney tower is *split* if each Whitney disk is embedded, and the set of singularities in the interior of any framed Whitney disk consists of either a single transverse intersection point, or a single boundary arc of a higher-order Whitney disk, or is empty; and if each non-trivially twisted Whitney disk has no singularities in its interior, and has twisting equal to ± 1. This can always be arranged by performing (twisted) finger moves along Whitney disks guided by arcs connecting the Whitney disk boundary arcs (see [10, Section 2.5]).

Splitting simplifies the combinatorics of Whitney tower constructions and will be assumed, often without mention, in subsequent sections. Splitting an order n (twisted) Whitney tower $\mathcal{W}$ does not change $\tau_n^\infty(\mathcal{W}) \in \mathcal{T}_n^\infty$ [10, Lemma 2.12].

2.9. *Intersection forests for split twisted Whitney towers*

Recall from [10, Definition 2.11] that the disjoint union of signed trees and ∞-trees associated to the unpaired intersections and ± 1-twisted Whitney disks in a split twisted Whitney tower $\mathcal{W}$ is denoted by $t(\mathcal{W})$, and called the *intersection forest* of $\mathcal{W}$. Here each tree t_p associated to an unpaired intersection p is equipped with the sign of p, and each ∞-tree J^∞ associated to a clean ± 1-twisted Whitney disk is given the corresponding sign ± 1.

In any split $\mathcal{W}$, the intersection forest can be thought of as an embedding of the disjoint union of trees $t(\mathcal{W})$ into $\mathcal{W}$ which embodies both the geometric and algebraic data associated to $\mathcal{W}$: If we think of the trees as subsets of $\mathcal{W}$, then all singularities of $\mathcal{W}$ are contained in a neighborhood of $t(\mathcal{W})$; and if we think of the trees as generators, then $t(\mathcal{W})$ is an 'abelian word' representing $\tau_n^\infty(\mathcal{W})$. (In any $\mathcal{W}$ of order n, it is always possible to eliminate all intersections of order strictly greater than n, for instance by performing finger moves ('pushing down') to create algebraically canceling pairs of order n intersections, see discussion in [10, Section 4].)

REMARK 28. In the older papers [8, 45, 47], 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 $t(\mathcal{W})$.

3. *Twisted Whitney towers and gropes*

For use in subsequent sections, this section recalls the correspondence between (split) Whitney towers and (dyadic) capped gropes [8, 45] in the 4-ball, and extends this relationship to the twisted setting. The main goal is to describe how this correspondence preserves the associated disjoint unions of signed trees (intersection forests). In particular, Lemma 31 will be used in Section 4 to exhibit the relationship between twisted Whitney towers and Milnor invariants in the proof of Theorem 6. A detailed understanding of the material in this section relies heavily on having digested (the proof of) [45, Theorem 5]. An illustration of the tree-preserving Whitney tower-grope correspondence can be seen in Figure 8.

3.1. *Dyadic gropes and their associated trees*

This subsection reviews and fixes some basic grope terminology. It will suffice to work with *dyadic* gropes, that is, gropes whose higher stages are all genus 1; these correspond to split Whitney towers (Subsection 2.8), and gropes in 4-manifolds can always be modified to be dyadic by Krushkal's 'grope splitting' operation [29].

A *dyadic grope* G is a 2-complex constructed by the following method.

(1) Start with a compact orientable connected surface of any positive genus, called the *bottom stage* of G, and choose a symplectic basis of circles on this bottom stage surface.

(2) Attach punctured tori to any number of the basis circles and choose hyperbolic pairs of circles on each attached torus.

(3) Iterate the second step a finite number of times, that is, attach punctured tori to any number of previously chosen basis circles that do not already have a torus attached to them, and choose hyperbolic circle-pairs for these new tori.

The attached tori are the *higher stages* of G, and at each iteration in the construction tori can be attached to circles in any stage. The basis circles in all stages of G that do not have a torus attached to them are called the *tips* of G.

Our requirement that the bottom stage of G has positive genus serves only to simplify notation and terminology, as the genus zero case will not be needed in our constructions.

Attaching 2-disks along all the tips of G yields a *capped* (dyadic) grope, denoted by G^c, and the uncapped grope G is called the *body* of G^c.

Cutting the bottom stage of G into genus 1 pieces decomposes G (and G^c) into *branches*, and our notion of dyadic grope (following [**8**, **45**]) is more precisely called a 'grope with dyadic branches' in [**29**].

With an eye toward defining intersection forests for capped gropes in B^4, we start by associating to an abstract capped grope G^c the following disjoint union $t(G^c)$ of unlabeled and unoriented unitrivalent trees: Assume first that the bottom stage of G^c is a genus 1 surface with boundary. Then define $t(G^c)$ to be the unitrivalent tree which is dual to the 2-complex G^c. Specifically, the tree $t(G^c)$ can be embedded in G^c in the following way. Choose a vertex in the interior of each surface stage and each cap of G^c. Then each edge of $t(G^c)$ is a sheet-changing path between vertices in adjacent stages or caps (here 'adjacent' means 'intersecting in a circle'). One univalent vertex of $t(G^c)$ sits in the bottom stage of G^c, each of the other univalent vertices is a point in the interior of a cap of G^c and each higher stage of G^c contains a single trivalent vertex of $t(G^c)$ (see, for example, the right-hand side of Figure 8).

In the case where the bottom stage of G^c has genus greater than 1, $t(G^c)$ is defined by cutting the bottom stage into genus 1 pieces and taking the disjoint union of the unitrivalent trees just described. Thus, each branch of G^c contains a single tree in $t(G^c)$.

Note that each tree in $t(G^c)$ has exactly one univalent vertex which sits in the bottom stage of G^c; these vertices can naturally be considered as roots, and it is customary to associate rooted trees to gropes. Here, we prefer to ignore this extra information, since we will be identifying $t(G^c)$ with the unrooted trees associated to Whitney towers.

The *class* of a capped grope G^c is one more than the minimum of the orders of the trees in $t(G^c)$. The body G of G^c inherits the same union of trees, $t(G) := t(G^c)$, and the same notion of class.

Convention: *For the rest of this paper, gropes may be assumed to be dyadic, even if not explicitly stated.*

3.2. *Intersection forests for capped gropes bounding links*

The *boundary* ∂G of a grope G is the boundary of its bottom stage. An embedding $(G, \partial G) \hookrightarrow (B^4, S^3)$ is *framed* if a disjoint parallel push-off of the bottom stage of G induces a given framing of $\partial G \subset S^3$ and extends to a disjoint parallel push-off of G in B^4.

DEFINITION 29. For a framed link $L \subset S^3$, the statement 'L bounds a capped grope G^c in B^4' means that the link components L_i bound disjointly embedded framed gropes $G_i \subset B^4$, such that the tips of the G_i bound framed caps whose interiors are disjointly embedded, with

each cap having a single transverse interior intersection with the bottom stage of some G_j. Here a cap is *framed* if the parallel push-off of its boundary in the grope extends to a disjoint parallel copy of the entire cap. The union of the gropes is denoted $G := \bigcup_i G_i$, and $G^c := \bigcup_i G_i^c$ is the union of G together with all the caps.

All previous grope notions carry over to this setting, even though the bottom stage of G is not connected; for example, we refer to the grope G as the body of the capped grope G^c. In particular, the disjoint union of trees $t(G^c) := \amalg_i t(G_i^c)$ can now be considered a subset of B^4. This provides labels from $\{1, 2, \ldots, m\}$ for all univalent vertices: The bottom-stage univalent vertex of each tree in $t(G_i^c)$ inherits the label i; and if a cap intersects the bottom stage of G_j, then the vertex corresponding to that cap inherits the label j, as shown in the right-hand side of Figure 8. Orientations on all stages of G induce orientations of the trivalent vertices in $t(G^c)$, and orientations on all caps determine signs for each cap-bottom stage intersection. To each tree in $t(G^c)$ is associated a sign $\pm$ which is the product of the signs of its caps. We assume the convention that the orientations of the bottom stages of G correspond to the link orientation. Thus, when G^c is oriented, meaning that all stages and caps are oriented, $t(G^c)$ is a disjoint union of signed oriented labeled trees which we call the *intersection forest* of G^c, in line with the terminology for Whitney towers.

3.3. *Intersection forests for twisted capped gropes bounding links*

A *twisted capped grope* G^c in B^4 is the same as a capped grope as in Definition 29 just above, except that at most one cap in each branch of G^c is allowed to be arbitrarily *twisted* as long as its interior is embedded and disjoint from all other caps and stages of G^c. Here a cap c is k-twisted, for $k \in \mathbb{Z}$, if the parallel push-off of its boundary in the grope determines a section of the normal bundle of $c \subset B^4$ with relative Euler number k. (So a 0-twisted cap is framed.)

DEFINITION 30. A link $L \subset S^3$ *bounds a twisted capped grope* if the link components L_i bound disjointly embedded framed gropes $G_i \subset B^4$ which extend to a twisted capped grope $G^c = \bigcup_i G_i^c \subset B^4$.

The *intersection forest* $t(G^c)$ *of a twisted capped grope bounding a link* is defined as the extension of the framed definition which labels each univalent vertex that corresponds to a non-trivially k-twisted cap with the twist symbol ∞, and takes the twisting k as a coefficient.

Recall from Subsection 3.1 that for a capped grope G^c, if n is the minimum of the orders of the trees in $t(G^c)$, then the class of G^c is $n+1$.

Motivated by the correspondence with twisted Whitney towers described below, we define the *class of a twisted capped grope* G^c to be $n+1$ if n is the minimum of the orders of the non-∞ trees in $t(G^c)$, and if the ∞-trees in $t(G^c)$ are of order at least $n/2$.

3.4. *From twisted Whitney towers to twisted capped gropes*

In [45], a 'tree-preserving' procedure for converting an order n (framed) Whitney tower $\mathcal{W}$ into a class $n+1$ capped grope (and vice versa) is described in detail. This construction will be extended to the twisted setting in Lemma 31 just below, which will be used in the proof of Theorem 6 given in the next section. The rough idea is that the 'subtower' of Whitney disks containing a tree in a split Whitney tower can be surgered to a dyadic branch of a capped grope containing the same tree, with the capped grope orientation inherited from that of the Whitney tower. Orientation and sign conventions will be presented during the course of the proof.

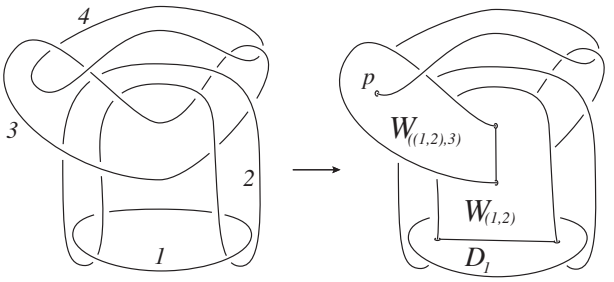

FIGURE 7. *Moving into B^4 from left to right, a Bing-doubled Hopf link $L \subset S^3$ bounds an order 2 Whitney tower $\mathcal{W}$: The order 0 disk D_1 consists of a collar on L_1 together with the indicated embedded disk on the right. The other three order 0 disks D_2, D_3 and D_4 consist of collars on the other link components which extend further into B^4 and are capped off by disjointly embedded disks. The Whitney disk $W_{(1,2)}$ pairs $D_1 \cap D_2$, and $W_{((1,2),3)}$ pairs $W_{(1,2)} \cap D_3$, with $p = W_{((1,2),3)} \cap D_4$, the only unpaired intersection point in $\mathcal{W}$.*

LEMMA 31. *If L bounds an order n split twisted Whitney tower $\mathcal{W}$, then L bounds a dyadic class $n+1$ twisted capped grope G^c such that*:

(i) *$t(\mathcal{W})$ is isomorphic to $t(G^c)$;*
(ii) *each framed cap of G^c has intersection $+1$ with a bottom stage of G, except that one framed cap in each dyadic branch of G^c with signed tree $\epsilon_p \cdot t_p$ has intersection ϵ_p with a bottom stage*;
(iii) *each ϵ-twisted cap of G^c contains the corresponding ∞-labeled vertex of its ∞-tree in $t(G^c)$.*

Proof. Outline: A detailed inductive proof of the framed unoriented case is given in [**45**, Theorem 5]. We will adapt the proof from [**45**] to the current twisted setting, sketching the construction while introducing orientation and sign conventions. The basic idea of the procedure is to tube (0-surger) along one boundary-arc of each Whitney disk; but in order to maximize the class of the resulting grope, Whitney moves may need to be performed when trees are not *simple*, meaning right- or left-normed (see [**45**, Figure 17 of Section 7]). As mentioned at the start of this section, an appreciation of the role played by these subtleties in the current proof depends largely on having understood the proof of [**45**, Theorem 5].

A simple example of the construction (in the framed case) is illustrated in Figures 7 and 8, which show how an order 2 Whitney tower bounded by the Bing-double of the Hopf link can be converted to a class 3 capped grope.

For each $t_p \in t(\mathcal{W})$, the construction works upward from a chosen Whitney disk having a boundary arc on an order 0 disk D_i, which corresponds to the choice of an i-labeled univalent vertex of t_p, creating caps out of Whitney disks, then turning these caps into surface stages whose caps are created from higher-order Whitney disks, and so on, until the process terminates when each framed cap has a single intersection with a bottom stage surface. The resulting dyadic branch of G^c will inherit the tree t_p as a subset. Similarly, for each ∞-tree $\pm J^\infty \in t(\mathcal{W})$, the construction will yield a dyadic branch containing J^∞ with the ∞-vertex sitting in a ± 1-twisted cap.

The surgery step: Figure 9 illustrates a surgery step and the corresponding modification of the embedded tree near a trivalent vertex corresponding to a Whitney disk $W_{(I,J)}$ in $\mathcal{W}$. The sheet c_I is a (temporary) cap which has already been created, or is just an order 0 disk D_i with $I = i$ in the first step of creating a dyadic branch of G^c. Any interior intersections of $W_{(I,J)}$ are not shown. After the surgery which turns the c_I into a surface stage S_I, the

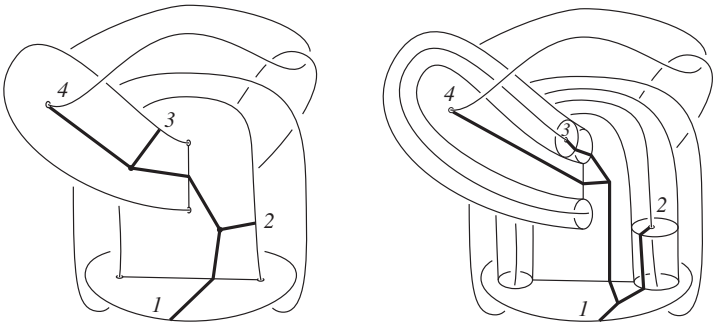

FIGURE 8. *Both sides of this figure correspond to the slice of B^4 shown in the right-hand side of Figure 7. The tree $t_p = \langle((1,2),3),4\rangle$ is shown on the left as a subset of the order 2 Whitney tower $\mathcal{W}$. Replacing this left picture by the picture on the right illustrates the tree-preserving construction of a class 3 capped grope G^c bounded by L. In this case, the component G^c_1 bounded by L_1 is the class 3 capped grope shown (partly translucent) on the right (together with a collar on L_1) which is gotten by surgering D_1 and $W_{(1,2)}$. The three other components of G are just the disks D_2, D_3 and D_4 of $\mathcal{W}$, each of which has a single intersection with a cap of G^c_1.*

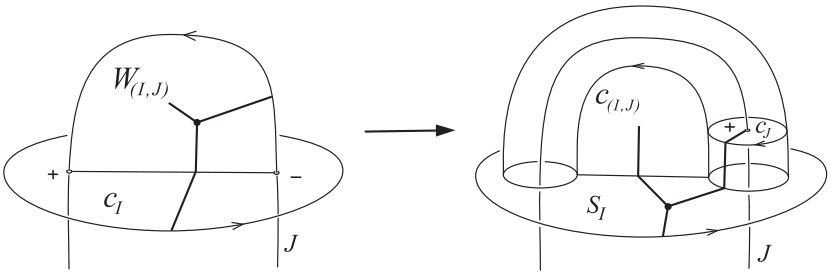

FIGURE 9. *A surgery step in the resolution of an order n twisted Whitney tower to a class $n+1$ twisted capped grope. Any interior intersections in $W_{(I,J)}$ and $c_{(I,J)}$ are not shown.*

Whitney disk $W_{(I,J)}$ minus part of a collar becomes one cap $c_{(I,J)}$ and a normal disk to the J-sheet becomes a dual cap c_J. The S_I stage inherits the orientation of c_I, and the cap $c_{(I,J)}$ inherits the orientation of $W_{(I,J)}$. As pictured in Figure 9, the effect of the surgery on the tree sends the trivalent vertex in $W_{(I,J)}$ to the trivalent vertex in the S_I sheet, with the induced orientation. The cap c_J is a parallel copy of what used to be a neighborhood in c_I around the negative intersection point paired by $W_{(I,J)}$, but with the opposite orientation, so that c_J has a single positive intersection with the J-sheet.

Here the I-subtree sits in the part of the grope branch which has already been constructed, while the J-subtree and any K-subtree corresponding to intersections with $c_{(I,J)}$ sit in subtowers of $\mathcal{W}$ which have yet to be converted into grope stages. The proof proceeds by considering the various cases depending on the orders of I, J and K.

We call a Whitney disk or cap *clean* if it is embedded and free of any interior intersections with any surface sheets.

The surgery cases: If $W_{(I,J)}$ had a single interior intersection with an order 0 disk D_k, then so does the cap $c_{(I,J)}$; and we relabel this cap c_k. If in this case $J = j$ is also order zero, then there is no further modification to c_j and c_k, which remain as caps intersecting the bottom stages of the gropes G_j and G_k when the construction is complete.

If $W_{(I,J)}$ was a clean ϵ-twisted Whitney disk, then $c_{(I,J)}$ is a clean ϵ-twisted cap of G^c containing the ω-label of the ω-tree associated to the branch. In this case, there is no further modification of the cap, which will be denoted by $c^{\omega}_{(I,J)}$.

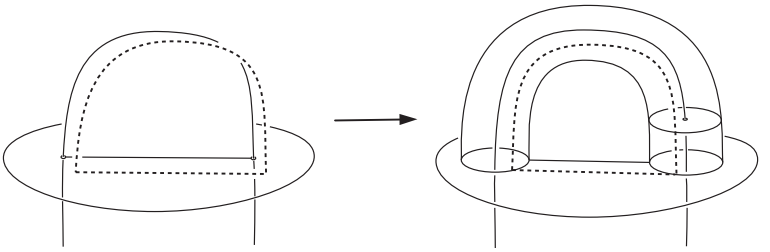

FIGURE 10. *The framing obstruction determined by the Whitney section over the boundary of a Whitney disk is passed on to the framing obstruction on the cap resulting from surgery.*

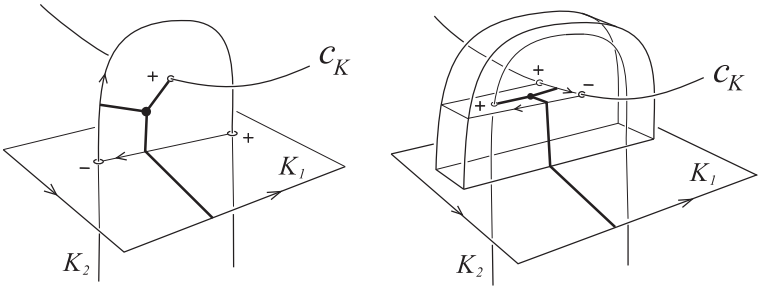

FIGURE 11. *A Whitney move preserves the sign and orientation at a trivalent vertex.*

Note that surgering Whitney disks to caps preserves twistings; see Figure 10.

If $J = j$ is order zero and $W_{(I,j)}$ was a clean ϵ-twisted Whitney disk yielding $c^{\infty}_{(I,j)}$, then there is also no further modification to the dual cap c_j. If $J = (J_1, J_2)$ has positive order, then the clean twisted cap $c^{\infty}_{(I,J)}$ remains 'as is', but the dual cap c_J is modified as described below (in the next-to-last paragraph of the proof).

If the cap $c_{(I,J)}$ contains an arc of a Whitney disk boundary, then the just-described surgery step for c_I applies to $c_{(I,J)}$. Otherwise, the grope construction requires a Whitney move as described next.

If the cap $c_{(I,J)}$ intersects some W_K transversely in the single point p, with $\text{sign}(p) = \epsilon_p$, and $K = (K_1, K_2)$ of positive order, then the grope construction proceeds by doing a Whitney move guided by W_K on either the K_1-sheet or the K_2-sheet: The effect of this W_K-Whitney move is to replace p by a Whitney-disk boundary-arc in $c_{(I,J)}$ so that the surgery step can be applied. Here p could be the original unpaired intersection in t_p, or an intersection created during the construction, and Figure 11 illustrates how the oriented tree and the sign of the unpaired intersection are preserved in the case $\epsilon_p = +1$; the case $\epsilon_p = -1$ can be checked in the same way.

Similarly, if $J = (J_1, J_2)$ has positive order, then the grope construction proceeds by doing a W_J-Whitney move to replace the positive intersection point between c_J and W_J by a boundary arc of a Whitney disk, so that the surgery step can be applied to c_J. That this preserves the oriented tree and the $+1$ sign of the unpaired intersection also follows from (a re-labeling of) Figure 11.

For each tree in $t(\mathcal{W})$, this procedure terminates when each framed cap has a single intersection with a bottom stage, creating a dyadic branch of the capped grope G^c; and applying the procedure to all trees in $t(\mathcal{W})$ yields G^c, containing its intersection forest $t(G^c)$, with all vertex orientations induced by the orientation of G^c. Since conditions (ii) and (iii) of the lemma

are satisfied, it follows that $t(G^c)$ and $t(\mathcal{W})$ are isomorphic, since the coefficients of the trees are also preserved. □

4. *Proof of Theorem* 6

Recall the content of Theorem 6: For L bounding an order n twisted Whitney tower $\mathcal{W}$, the first non-vanishing order n Milnor invariant $\mu_n(L)$ can be computed from $\mathcal{W}$ as

$$\mu_n(L) = \eta_n(\tau_n^{\infty}(\mathcal{W})),$$

where $\mu_n(L) := \sum_i X_i \otimes \mu_n^i(L) \in \mathsf{L}_1 \otimes \mathsf{L}_{n+1}$ collects the length $n+1$ iterated commutators determined by the link longitudes considered as Lie brackets $\mu_n^i(L)$ in the free $\mathbb{Z}$-Lie algebra, and the map η_n converts trees into rooted trees (Lie brackets) by summing over all choices of roots (the definition of η_n is recalled below). The proof of this statement will also show that $\mu_k(L)$ vanishes for all $k < n$.

We first note that Theorem 6 holds in two special cases.

The order 0 *case*: It is easily checked from the definitions in Subsection 1.1 that for $i \neq j$ the coefficient of $X_i \otimes X_j$ in $\mu_0(L)$ is the linking number of L_i and L_j, which via the well-known computation of linking numbers by counting signed intersections between the properly immersed disks D_i and D_j bounded by L_i and L_j is also equal to the coefficient of $X_i \otimes X_j$ in $\eta_0(\tau_0^{\infty}(\mathcal{W}))$.

Although Milnor invariants are not usually defined for knots, for framed links it is natural to consider the framing f_i of L_i as an order 0 (length 2) integer Milnor invariant, and the coefficient of $X_i \otimes X_i$ in $\mu_0(L)$ is exactly f_i when this framing is used to determine the ith longitude. To see that the coefficient in $\eta_0(\tau_0^{\infty}(\mathcal{W}))$ of $X_i \otimes X_i$ is also equal to f_i, let d_i denote the number of positive self-intersections of D_i minus the number of negative self-intersections of D_i. Then the relative Euler number of D_i with respect to the framing f_i on $L_i = \partial D_i$ is equal to $f_i - 2d_i$ (see, for example, [**10**, Figure 19] and accompanying discussion), and the terms of $\tau_0^{\infty}(\mathcal{W})$ which contribute via η_0 to the coefficient of $X_i \otimes X_i$ are exactly $(d_i) \cdot i \longrightarrow i + (f_i - 2d_i) \cdot \infty \longrightarrow i$, which get sent by η_0 to $(f_i) \cdot X_i \otimes X_i$.

So we may assume in the rest of the proof that $\mathcal{W}$ is of positive order. Note that this means that all link component framings are zero, since the self-intersections of all order 0 disks come in canceling pairs (paired by order 1 Whitney disks).

The case of slice links: In the case that $\mathcal{W}$ consists of disjointly embedded slice disks for L, Theorem 6 is true for all $n > 0$ since $\mu_n(L)$ vanishes, and so does $\tau_n^{\infty}(\mathcal{W})$ since $\mathcal{W}$ is an order n twisted Whitney tower for all n.

So in the rest of the proof we may also assume that the intersection forest $t(\mathcal{W})$ (defined in Subsection 2.9) is non-empty.

4.1. *Outline of the proof*

To prove Theorem 6, we will first convert the order n twisted Whitney tower $\mathcal{W}$ bounded by L to an order $n+1$ twisted capped grope G^c, as in Lemma 31. It will follow from an extension of *grope duality* [**30**] to the setting of twisted capped gropes, together with Dwyer's theorem [**15**], that we can compute the link longitudes in $\pi_1(B^4 \setminus G^c)$ instead of $\pi_1(S^3 \setminus L)$. Via the capped grope duality construction the iterated commutators determined by the longitudes will be seen to correspond exactly to the image of $\tau_n^{\infty}(\mathcal{W})$ under the map η_n. To preview the computation of the longitudes, the reader can examine Figures 7 and 8 which show the Whitney tower-to-capped grope conversion for L the Bing-double of the Hopf link. It should be clear from the right-hand side of Figure 8 that the longitude for component L_1 is a triple commutator $[x_2, [x_3, x_4]]$ of meridians to the other components, as exhibited by the class 3 capped grope G_1^c bounded by L_1 and containing the order 2 tree. As a consequence of grope duality, it will turn

out that the other longitudes also bound class 3 gropes which correspond to choosing roots at the other univalent vertices on the same order 2 tree (although these gropes are not so visible in the figure).

For the reader's convenience, we recall from the introduction the definition of the map $\eta_n : \mathcal{T}_n^{\infty} \to \mathsf{D}_n$, where $\mathsf{D}_n = \mathsf{D}_n(m)$ is the kernel of the bracket map $\mathsf{L}_1 \otimes \mathsf{L}_{n+1} \to \mathsf{L}_{n+2}$.

For v a univalent vertex of an order n tree t, denote by $B_v(t) \in \mathsf{L}_{n+1}$ the Lie bracket of generators $X_1, X_2, \ldots, X_m$ determined by the formal bracketing from $\{1, 2, \ldots, m\}$ which is gotten by considering v to be a root of t.

Denoting the label of a univalent vertex v by $\ell(v) \in \{1, 2, \ldots, m\}$, the map $\eta_n : \mathcal{T}_n^{\infty} \to \mathsf{L}_1 \otimes \mathsf{L}_{n+1}$ is defined on generators by

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

where the first sum is over all univalent vertices v of t, and the second expression is indeed in $\mathsf{L}_1 \otimes \mathsf{L}_{n+1}$ since the coefficient of $\eta_n(\langle J, J \rangle)$ is even.

The following lemma is proved in Subsection 4.3.

LEMMA 32. *The homomorphism $\eta_n : \mathcal{T}_n^{\infty} \to \mathsf{D}_n$ is a well-defined surjection.*

We will also use the following lemma.

LEMMA 33. *If $L \subset S^3$ bounds a class $(n+1)$ twisted capped grope $G^c \subset B^4$, then the inclusion $S^3 \setminus L \hookrightarrow B^4 \setminus G^c$ induces an isomorphism*

$$\frac{\pi_1(S^3 \setminus L)}{\pi_1(S^3 \setminus L)_{n+2}} \cong \frac{\pi_1(B^4 \setminus G^c)}{\pi_1(B^4 \setminus G^c)_{n+2}}.$$

The proof of Lemma 33 is given below in Subsection 4.4.

4.2. *Computing the longitudes*

By Lemma 33, we can compute the iterated commutators determined by the link longitudes in $\pi_1(B^4 \setminus G^c)$ modulo $\pi_1(B^4 \setminus G^c)_{n+2}$. The computation will show that the longitudes lie in $\pi_1(B^4 \setminus G^c)_{n+1}$, which implies that $\mu_k(L)$ vanishes for all $k < n$.

Terminology note: Throughout this proof, we will use the word *meridian* to refer to fundamental group elements represented by normal circles to deleted surfaces in 4-space; and on occasion such circles will themselves be referred to as 'meridians'.

Conventions: Via the isomorphisms of Lemma 33 and Subsection 1.1, we make the identifications

$$\frac{\pi_1(B^4 \setminus G^c)_{n+1}}{\pi_1(B^4 \setminus G^c)_{n+2}} \cong \frac{\pi_1(S^3 \setminus L)_{n+1}}{\pi_1(S^3 \setminus L)_{n+2}} \cong \frac{F_{n+1}}{F_{n+2}},$$

where the generators $\{x_1, x_2, \ldots, x_m\}$ are meridians to the bottom stages of G^c, with x_i chosen to have linking number $+1$ with the bottom stage of the grope component G_i which is bounded by L_i.

Orientations of surface sheets and their boundary circles are related by the usual 'outward vector first' convention.

We use the commutator notation $[g, h] := ghg^{-1}h^{-1}$ and exponential notation $g^h := hgh^{-1}$ for group elements g and h.

Since an element in F_{n+1}/F_{n+2} determined by an $(n+1)$-fold commutator of elements of F/F_{n+2} only depends on the conjugacy classes of the elements, we can and will suppress

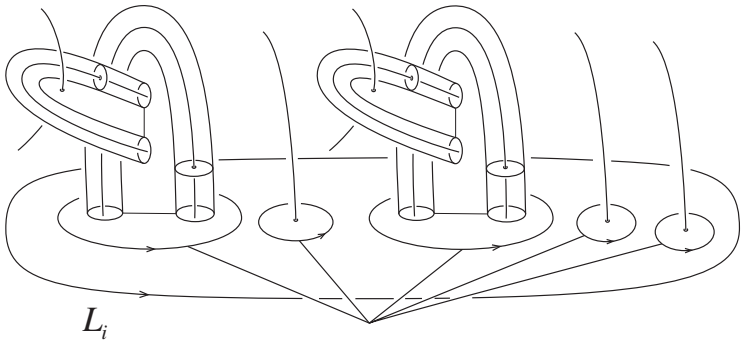

FIGURE 12. *A parallel push-off of L_i is isotopic to a product of loops which are boundary circles of parallel push-offs of dyadic branches of G_i or meridional circles to framed caps of G^c. So the corresponding factors γ_{i_r} of the ith longitude $\gamma_i = \prod_r \gamma_{i_r}$ are in one-to-one correspondence with the i-labeled vertices of the trees in $t(G^c)$.*

basings of meridians from computations. This follows easily from the commutator relation $[xy, z] = [y, z]^x[x, z]$ which holds in any group. The following relations in F_{n+1}/F_{n+2} will be useful.

For any length $n+1$ commutator $[x_I, x_J]$, and $\epsilon = \pm 1$,

$$[x_I, x_J^\epsilon] = [x_I^\epsilon, x_J] = [x_J^{-\epsilon}, x_I] = [x_J^\epsilon, x_I]^{-1} = [x_J^\epsilon, x_I^{-1}]. \tag{4.1}$$

For L bounding $\mathcal{W}$ of positive order, the longitudes γ_i are represented by 0-parallel push-offs of the link components. As illustrated in Figure 12, each longitude factors as $\gamma_i = \prod_r \gamma_{i_r}$, with each γ_{i_r} represented by either a parallel push-off of the boundary of a dyadic branch of G_i or a meridian to a framed cap in G^c. (The ordering of the factors of γ_i is irrelevant since F_{n+1}/F_{n+2} is abelian.)

For each i, the factors γ_{i_r} are in one-to-one correspondence with the set of i-labeled vertices v_{i_r} on all the trees in $t(G^c)$ (since each i-labeled univalent vertex on a tree corresponds either to an intersection between a framed cap and the bottom stage of G_i or to the i-labeled vertex sitting in the bottom stage of a dyadic branch of G_i). To complete the proof of Theorem 6, it suffices to check that each γ_{i_r} is equal to the iterated commutator $\beta_{v_{i_r}}(t)^\epsilon \in F_{n+1}/F_{n+2}$ determined by putting a root at v_{i_r} on the tree $\epsilon \cdot t \in t(G^c)$ containing v_{i_r} or equal to $\beta_{v_{i_r}}(\langle J, J\rangle)^\epsilon$ for v_{i_r} in an ∞-tree $\epsilon \cdot J^\infty \in t(G^c)$. Here the correspondence between rooted trees and iterated commutators is directly analogous to the correspondence with Lie brackets, and the isomorphism $F_{n+1}/F_{n+2} \cong \mathsf{L}_{n+1}$ in the definition (1.1) of $\mu_n(L)$ maps $\beta_{v_{i_r}}(t)^\epsilon$ to the Lie bracket $\epsilon \cdot B_{v_{i_r}}(t) \in \mathsf{L}_{n+1}$ as in the definition of the map η_n. Similarly for ∞-trees J^∞, $\beta_{v_{i_r}}(\langle J, J\rangle)^\epsilon$ maps to the correct Lie bracket $\epsilon \cdot B_{v_{i_r}}(\langle J, J\rangle) \in \mathsf{L}_{n+1}$ if n is even. Then $\mu_n^i(L)$ is the sum of these Lie brackets over all v_{i_r}.

4.2.1. *The order* 1 *case.* As a warm-up and base case for the general proof, we check that η_1 takes $\tau_1^\infty(\mathcal{W})$ to $\mu_1(L)$ (the 'triple linking numbers' of L) for any order 1 twisted Whitney tower $\mathcal{W}$ bounded by L. In this case, the grope construction yields a class 2 twisted capped grope G^c bounded by L, with intersection forest $t(G^c)$ a disjoint union of signed order 1 Y-trees representing $\tau_1^\infty(\mathcal{W})$. The body G is just a collection of disjointly embedded surfaces, and there are no twisted caps (since odd-order twisted Whitney towers do not contain twisted Whitney disks).

First consider the case where $t(G^c) = \epsilon_p \cdot t_p = \epsilon_p \cdot \langle (i, j), k\rangle$ is a single Y-tree, with i, j and k distinct, and G consists of a single genus 1 surface G_i bounded by L_i, together with disjointly embedded disks G_j and G_k bounded by the link components L_j and L_k (Figure 13).

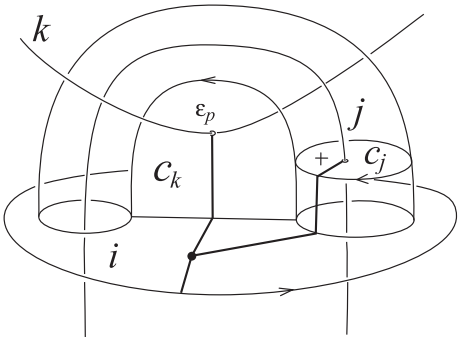

FIGURE 13. *Near the trivalent vertex of the signed* Y*-tree* $\epsilon_p \cdot t_p = \epsilon_p \cdot \langle (i,j),k\rangle$ *in a dyadic class* 2 *capped grope component* (*capped surface*) *bounded by* L_i.

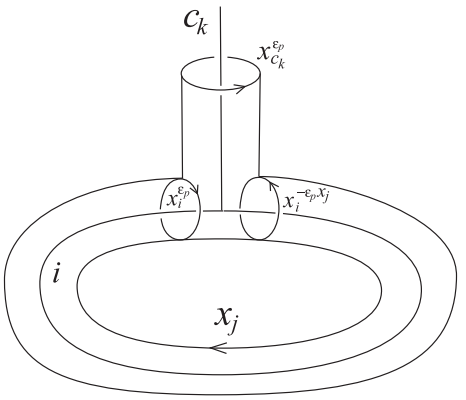

FIGURE 14. *A meridian to the cap* c_k *in Figure* 13 *bounds a genus* 1 *surface which is a punctured normal torus to the surface stage containing the cap boundary. This normal torus consists of circle fibers in the normal circle bundle over a dual circle to the cap boundary in the surface stage. This dual circle is parallel to the boundary of the dual cap* (*which in Figure* 13 *represents the meridian* x_j). *Since the* (*closed*) *normal torus has a single intersection with the cap, it is also called a 'dual torus' for the cap.*

We want to check that:

$$\mu_1(L) = \eta_1(\epsilon_p \cdot \langle (i,j),k\rangle) = \epsilon_p \cdot X_i \otimes \prec {}^{k}_{j} + \epsilon_p \cdot X_j \otimes {}_i \prec {}^{k} + \epsilon_p \cdot X_k \otimes {}_i \prec {}_{j}\,.$$

A parallel push-off of L_i bounds a parallel push-off of G_i in $B^4 \setminus G^c$ and the longitude γ_i can be computed from Figure 13 (using the commutator relations (4.1)):

$$\gamma_i = [x_j^{-1}, x_k^{-\epsilon_p}] = [x_j, x_k]^{\epsilon_p}.$$

This is the correct commutator $\beta_{v_i}(t_p)^{\epsilon_p} \in F_2/F_3$ corresponding to choosing a root for t_p at the i-labeled vertex v_i, confirming the first term in the right-hand side of the above expression for $\mu_1(L)$:

$$X_i \otimes \epsilon_p \cdot B_{v_i}(t_p) = X_i \otimes \epsilon_p \cdot [X_j, X_k] = X_i \otimes \mu_1^i(L).$$

A parallel push-off of L_k bounds a parallel push-off of the embedded disk G_k in $B^4 \setminus G$, with G_k intersecting G^c in the single point $p \in c_k$ with sign ϵ_p. Thus, the longitude γ_k is equal to $x_{c_k}^{\epsilon_p} \in F/F_{n+2}$, the positive meridian x_{c_k} to the cap c_k raised to the power ϵ_p. This meridian can be expressed in terms of the generators using the 'dual torus' to c_k illustrated in

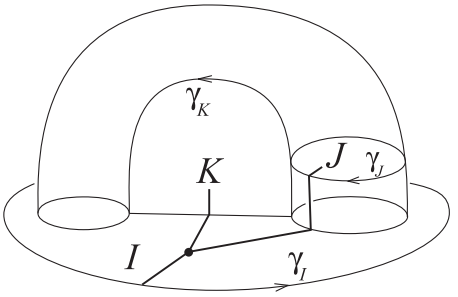

FIGURE 15. *Near a trivalent vertex in a dyadic branch of* G^c.

Figure 14, giving:

$$\gamma_k = x_{c_k}^{\epsilon_p} = x_i^{\epsilon_p} x_i^{-\epsilon_p x_j} = x_i^{\epsilon_p} x_j x_i^{-\epsilon_p} x_j^{-1} = [x_i^{\epsilon_p}, x_j] = [x_i, x_j]^{\epsilon_p},$$

which is the correct commutator $\beta_{v_k}(t_p)^{\epsilon_p}$ when the root of t_p is at the k-labeled vertex v_k. (One way to check this expression for $x_{c_k}^{\epsilon_p}$ directly from Figure 13 is to push the k-sheet down off c_k into the i-sheet by a finger move (the vertical tube in Figure 14) to get a canceling pair of intersection points which correspond to the factors $x_i^{\epsilon_p}$ and $x_i^{-\epsilon_p x_j}$.) This confirms the third term in the right-hand side of the expression for $\mu_1(L)$:

$$X_k \otimes \epsilon_p \cdot B_{v_k}(t_p) = X_k \otimes \epsilon_p \cdot [X_i, X_j] = X_k \otimes \mu_1^k(L).$$

By similarly using a dual torus, one can also check that the contribution to γ_j coming from the intersection point in the cap c_j is equal to $\beta_{v_j}(t_p)^{\epsilon_p} = [x_k, x_i]^{\epsilon_p}$, confirming the second term in the right-hand side of the expression for $\mu_1(L)$:

$$X_j \otimes \epsilon_p \cdot B_{v_j}(t_p) = X_j \otimes \epsilon_p \cdot [X_k, X_i] = X_j \otimes \mu_1^j(L).$$

Since all other link components bound disjointly embedded disks, this confirms Theorem 6 in this case where $t(\mathcal{W}) = t(G^c) = \epsilon_p \cdot t_p = \epsilon_p \cdot \langle (i,j), k \rangle$ with i, j and k distinct. If i, j and k are not distinct, then $t_p = 0 \in \mathcal{T}_1^\infty$ by the boundary-twist relations $\langle (i,j), j \rangle = 0$ (and the just-described computation will show that t_p contributes trivially to $\mu_1(L)$, since $[x_j, x_j] = 0$ and $[x_j, x_i] + [x_i, x_j] = 0$). The general order 1 case follows by summing the above computation over all factors of each longitude.

4.2.2. *The higher-order framed case.* Now consider the general order n case with the assumption that $\mathcal{W}$ contains no twisted Whitney disks, so that G^c is a class $n+1$ capped grope with no twisted caps. That the longitude factors are equal to the iterated commutators corresponding to putting roots at the univalent vertices of $t(G^c)$ for $n > 1$ will follow by applying the computations for $n = 1$ to recursively express the relations between meridians and push-offs of boundaries of surface stages of G^c at an arbitrary trivalent vertex of $t(G^c)$. As before, we start by considering the case where G^c consists of a single dyadic branch containing $t_p = t(G^c)$.

Figure 15 shows three surface stages in G^c around a trivalent vertex which decomposes the (un-rooted) tree t_p into three (rooted) subtrees I, J and K (whose roots are identified at the trivalent vertex), with the I-subtree reaching down to the bottom stage of G^c, and where we assume for the moment that J and K are of positive order (so the J- and K-sheets are not caps). Push-offs of the boundaries of the stages represent fundamental group elements γ_I, γ_J and γ_K; and we denote by x_I, x_J and x_K meridians to these stages.

The same computations as in the $n = 1$ case now give the three relations:

$$\gamma_I = [\gamma_J, \gamma_K], \quad x_J = [\gamma_K, x_I] \text{ and } x_K = [x_I, \gamma_J].$$

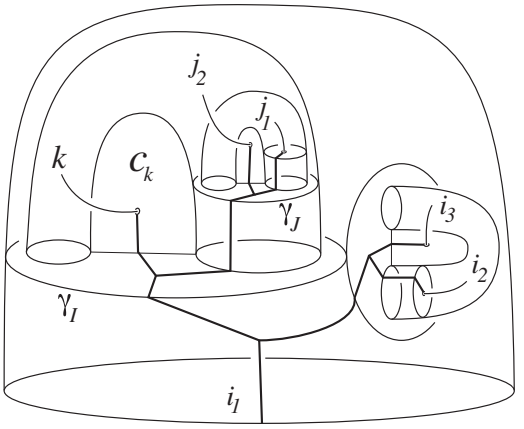

FIGURE 16. *An example of Figure 15 with $I = (i_1, (i_2, i_3))$ of order 2, $J = (j_1, j_2)$ of order 1 and $K = k$ of order zero.*

If either of J or K is order zero, say $K = k$, then the corresponding cap c_k intersects the bottom stage G_k, and so the cap boundary (labeled γ_K in Figure 15) will be a meridian x_k to G_k, and the cap meridian x_K will be denoted by x_{c_k}; and the relations become:

$$\gamma_I = [\gamma_J, x_k], \quad x_J = [x_k, x_I] \text{ and } x_{c_k} = [x_I, \gamma_J].$$

It follows recursively that, when J and K are of positive order, each of γ_I, x_J and x_K are equal to the iterated commutators in the generators corresponding to I, J and K:

$$\gamma_I = [J, K], \quad x_J = [K, I] \text{ and } x_K = [I, J].$$

And if $K = k$ is order zero, then we have

$$\gamma_I = [J, x_k], \quad x_J = [x_k, I] \text{ and } x_{c_k} = [I, J],$$

with similar relations for order zero $J = j$.

Theorem 6 is confirmed in this case by taking any of I, J and K to be order zero, which shows that the corresponding factor contributing to the longitude is the iterated commutator gotten by putting a root at that univalent vertex on t_p. The general framed case follows by summing this computation over all dyadic branches.

For instance, referring to the example of Figure 16 in the case $n = 4$, the contribution to the longitude γ_k coming from the pictured intersection between G_k and the cap c_k is represented by the cap meridian:

$$x_{c_k} = [x_I, \gamma_J] = [[x_{i_1}, [x_{i_2}, x_{i_3}]], [x_{j_1}, x_{j_2}]],$$

which is the iterated commutator $\beta_{v_k}(\langle (I, J), k\rangle)$ determined by putting a root at the k-labeled vertex v_k of the tree $\langle (I, J), k\rangle$.

4.2.3. *The general twisted case.* Now consider the general order n case where G^c may contain twisted caps (for even n) corresponding to ± 1-twisted Whitney disks (of order $n/2$) in $\mathcal{W}$. Again, by additivity of the computation over the dyadic branches, it is enough to consider a single dyadic branch of G^c containing a ± 1-twisted cap c_J^∞ and check that the corresponding ∞-tree J^∞ contributes $\eta_n(J^\infty) = \frac{1}{2}\eta_n(\langle J, J\rangle)$ to $\mu_n(L)$.

The key observation in this case is that because the cap c_J^∞ is ± 1-twisted, the element γ_∞ represented by a parallel push-off of the (oriented) boundary of the cap is the $(\pm)$-meridian $x_J^{\pm 1}$ to the cap. For $J = (J_1, J_2)$, referring to Figure 17 and using the same dual torus as for a

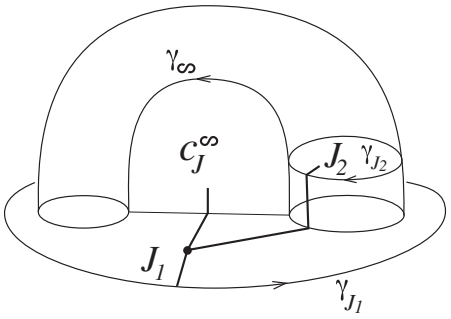

FIGURE 17. *Near a twisted cap in a dyadic branch of G^c.*

framed cap (Figure 14), this element can be expressed as the commutator:

$$\gamma_\infty = x_{(J_1,J_2)}^{\pm 1} = [x_{J_1}, \gamma_{J_2}]^{\pm 1},$$

where if $J_2 = j_2$ is order zero, then γ_{J_2} is replaced by the meridian x_{j_2} to G_{j_2} (as in the notation for the previous untwisted case).

So the analogous computations as in Figure 15 applied to the twisted setting of Figure 17 give the relations:

$$\gamma_{J_1} = [\gamma_{J_2}, [x_{J_1}, \gamma_{J_2}]]^{\pm 1} \quad \text{and} \quad x_{J_2} = [[x_{J_1}, \gamma_{J_2}], x_{J_1}]^{\pm 1},$$

and recursively as in the framed case:

$$\gamma_{J_1} = [J_2, [J_1, J_2]]^{\pm 1} \quad \text{and} \quad x_{J_2} = [[J_1, J_2], J_1]^{\pm 1},$$

with J_1 and J_2 denoting the corresponding iterated commutators in the meridional generators.

To see that the contribution to γ_{i_r} corresponding to any i_r-labeled vertex v_{i_r} of J^∞ is the iterated commutator $\beta_{v_{i_r}}(\langle J, J\rangle)$, observe that if v_{i_r} is in J_2, then the contribution will be an iterated commutator containing x_{J_2}, and if v_{i_r} is in J_1, then the contribution will be the iterated commutator containing γ_{J_1}. Thus, the effect of the twisted cap is to 'reflect' the iterated commutator determined by J at the ∞-labeled root. For instance, in the example of Figure 18 for the case $n = 8$, the contribution to the longitude γ_{j_1} corresponding to the boundary of the dyadic branch is:

$$\begin{aligned}
[[x_{j_2}, x_{j_3}], \gamma_{J_1}] &= [[x_{j_2}, x_{j_3}], [\gamma_{J_2}, \gamma_\infty]] \\
&= [[x_{j_2}, x_{j_3}], [[x_{j_4}, x_{j_5}], x_{c_J}]] \\
&= [[x_{j_2}, x_{j_3}], [[x_{j_4}, x_{j_5}], [x_{J_1}, \gamma_{J_2}]]] \\
&= [[x_{j_2}, x_{j_3}], [[x_{j_4}, x_{j_5}], [[x_{j_1}, [x_{j_2}, x_{j_3}]], [x_{j_4}, x_{j_5}]]]] \\
&= \beta_{v_{j_1}}(\langle J, J\rangle),
\end{aligned}$$

for $J = (J_1, J_2) = ((j_1, (j_2, j_3)), (j_4, j_5))$ and assuming that the twisting of c_J^∞ is $+1$.

Since each univalent vertex of J contributes one term to $\mu_n(L)$, the total contribution of the branch is equal to $\eta_n(J^\infty) = \frac{1}{2}\eta_n(\langle J, J\rangle)$.

This completes the proof of Theorem 6, modulo the proofs of Lemmas 32 and 33, which follow.

4.3. *Proof of Lemma* 32

Lemma 32 states that $\eta_n : \mathcal{T}^\infty \to \mathsf{D}_n$ is a well-defined surjection. Levine showed in [**32**] that an analogous map $\eta_n' : \mathcal{T}_n \to \mathsf{D}_n' := \mathrm{Ker}\{\mathsf{L}_1' \otimes \mathsf{L}_{n+1}' \to \mathsf{L}_{n+2}'\}$ is a well-defined surjection, where η_n'

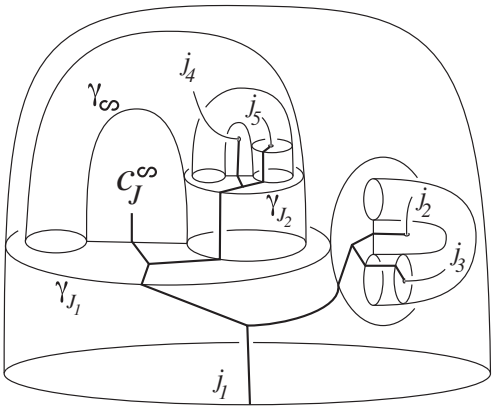

FIGURE 18. *An example of Figure 17 with* $J_1 = (j_1, (j_2, j_3))$ *of order* 2, *and* $J_2 = (j_4, j_5)$ *of order* 1.

is defined on trees using the same 'sum over all choices of root' formula as η_n, and $\mathsf{L}' = \bigoplus_{n\in\mathbb{N}} \mathsf{L}'_n$ is the free *quasi Lie algebra* on $\{X_1, X_2, \ldots, X_m\}$ gotten from the free $\mathbb{Z}$-Lie algebra L by replacing the self-annihlation relations $[X, X] = 0 \in \mathsf{L}$ with the AS relations $[X, Y] + [Y, X] = 0 \in \mathsf{L}'$. (See also [11], where we prove the *Levine Conjecture*: that η'_n is an isomorphism for all n.) It follows that η_n vanishes on the usual IHX and AS relations, and maps onto D_n. So it suffices to check that η_n respects the other relations in $\mathcal{T}_n^\infty$.

First consider the odd order case. To see that η_n vanishes on the boundary-twist relations, observe that $\eta_n(\langle\langle(i, J), J\rangle) = 0$, since placing a root at the i-labeled vertex determines a trivial symmetric bracket in L_{n+1}, and all the other Lie brackets come in canceling pairs corresponding to putting roots on vertices in each of the isomorphic J sub-trees.

Now considering the even-order case, we have

$$\eta_n((-J)^\infty) = \tfrac{1}{2}\eta_n(\langle -J, -J\rangle) = \tfrac{1}{2}\eta_n(\langle J, J\rangle) = \eta_n(J^\infty)$$

and

$$\eta_n(2J^\infty) = 2 \cdot \eta_n(J^\infty) = 2 \cdot \tfrac{1}{2}\eta_n(\langle J, J\rangle) = \eta_n(\langle J, J\rangle).$$

And for I, H and X, the terms in a Jacobi relation $I = H - X$, we have

$$\begin{aligned}
\eta_n(I^\infty) &= \frac{1}{2}\eta_n(\langle I, I\rangle) \\
&= \frac{1}{2}\eta_n(\langle H - X, H - X\rangle) \\
&= \frac{1}{2}(\eta_n(\langle H, H\rangle) - 2 \cdot \eta_n(\langle H, X\rangle) + \eta_n(\langle X, X\rangle)) \\
&= \eta_n(H^\infty + X^\infty - \langle H, X\rangle),
\end{aligned}$$

where the second equality comes from applying the Jacobi relation to the sub-trees I inside the inner product $\langle I, I\rangle$ and expanding. □

4.4. *Proof of Lemma* 33

By Dwyer's theorem [15], it suffices to show that the inclusion $(S^3 \setminus L) \hookrightarrow (B^4 \setminus G^c)$ induces an isomorphism on first homology and that the relative (integral) second homology group $H_2(B^4 \setminus G^c, S^3 \setminus L)$ is generated by maps of closed gropes of class at least $n + 2$ (where a grope is *closed* if its bottom stage is compact with empty boundary).

Observe first that $H_1(S^3 \setminus L)$ is Alexander dual to $H^1(L)$ and is hence generated by meridians. Similarly, $H_1(B^4 \setminus G^c)$ is generated by meridians to the bottom stages of the grope. It follows that the inclusion induces an isomorphism on H_1, since meridians of the link go to meridians of the bottom stages.

By Alexander duality, the generators of $H_2(B^4 \setminus G^c)$ which do not come from the boundary are the Clifford tori (or 'linking tori', see, for example, [**18**, 1.1, 2.1]) around the intersections between the caps and the bottom stages of G^c. Each such Clifford torus contains a pair of dual circles, one a meridian x_k to the kth bottom stage of G and the other a meridian x_{c_k} to the cap c_k. Referring to Figures 13 and 14 (with i and j replaced by I and J, respectively), the cap meridian x_{c_k} bounds a genus 1 surface T_{c_k} containing a pair of dual circles, one a meridian x_I to the I-stage of G^c containing ∂c_k and the other a parallel push-off of the boundary of the J-stage representing γ_J which is dual to c_k (if c_k is dual to another cap c_j, then γ_J is just a meridian x_j to the jth bottom stage of G^c).

Consider first the case where the dyadic branch of G^c containing c_k does not contain a twisted cap, and let $t = \langle k, (I, J)\rangle \in t(G^c)$ be the corresponding order n tree. Applying the grope duality construction of [**30**, Section 4] to T_{c_k} yields a class $n+1$ grope in $B^4 \setminus G$ having T_{c_k} as a bottom stage and associated tree $\langle k, (I, J)\rangle$ (with the k-labeled univalent vertex corresponding to T_{c_k}). Since this class $n+1$ grope consists of normal tori and parallel push-offs of higher stages of G, it actually lies in the complement of the caps of G^c (which only intersect the bottom stages of G). The union of this class $n+1$ grope with the Clifford torus is a class $n+2$ closed grope with associated order $n+1$ rooted tree $(k, (I, J))$ (with the root corresponding to the Clifford torus), completing the proof in the case where G^c has no twisted caps.

Now consider the case where the dyadic branch of G^c containing c_k does contain a twisted cap c^∞_J, with associated ∞-tree J^∞. Recall the observation of Paragraph 4.2.3 that a normal push-off of the cap boundary ∂c^∞_J representing $\gamma_\infty \in \pi_1(B^4 \setminus G^c)$ is a meridian x_J to c^∞_J. In this case, the grope duality construction of the previous paragraph which builds a grope on T_{c_k} will at some step look for a subgrope bounded by a normal push-off of ∂c^∞_J. Just as the computations in Subsection 4.2 show that x_J represents the iterated commutator in $\pi_1(B^4 \setminus G^c)$ corresponding to the rooted tree J, the punctured dual torus to c^∞_J bounded by x_J extends to a grope in $B^4 \setminus G^c$ with tree J. Thus, the torus T_{c_k} extends to (a map of) a grope in $B^4 \setminus G^c$ whose associated tree is gotten by putting a root at the corresponding k-labeled univalent vertex of (either one of the sub-trees) J in $\langle J, J\rangle$. It follows that the Clifford torus near the cap c_k extends to a grope whose corresponding tree is gotten by inserting an (rooted) edge into the edge of $\langle J, J\rangle$ adjacent to the k-labeled univalent vertex. Since the order of $\langle J, J\rangle$ is n, it follows that the class of the grope containing the Clifford torus as a bottom stage is $n+2$. □

5. *The order* 2 *twisted intersection invariant and the classical Arf invariant*

This section contains a proof of Lemma 10 from the introduction.

Recall the statements of Lemma 10: Any knot K bounds a twisted Whitney tower $\mathcal{W}$ of order 2 and the classical Arf invariant of K can be identified with the intersection invariant $\tau^\infty_2(\mathcal{W}) \in \mathcal{T}^\infty_2(1) \cong \mathbb{Z}_2$; and more generally, the classical Arf invariants of the components of an m-component link give an isomorphism $\mathrm{Arf} : \mathrm{Ker}(\mu_2 : \mathsf{W}^\infty_2 \twoheadrightarrow \mathsf{D}_2) \xrightarrow{\cong} (\mathbb{Z}_2 \otimes \mathsf{L}_1) \cong (\mathbb{Z}_2)^m$.

Proof. Starting with the first statement, observe that any knot $K \subset S^3$ bounds an immersed disk $D \looparrowright B^4$, and by performing cusp homotopies as needed it can be arranged that all self-intersections of D come in canceling pairs admitting order 1 Whitney disks. These Whitney disks can be made to have disjointly embedded boundaries by a regular homotopy applied to Whitney disk collars [**46**, Figure 3]. It is known that the sum modulo 2 of the number of

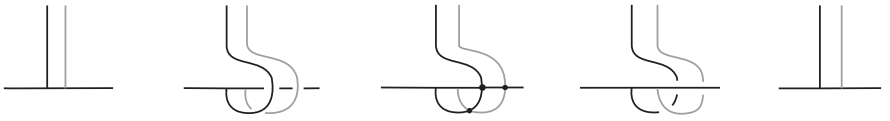

FIGURE 19. *Boundary-twisting a Whitney disk W changes $\omega(W)$ by ± 1 and creates an intersection point with one of the sheets paired by W. The horizontal arcs trace out part of the sheet, the dark non-horizontal arcs trace out the newly twisted part of a collar of W and the gray arcs indicate part of the Whitney section over W. The bottom-most intersection in the middle picture corresponds to the ± 1-twisting created by the move.*

intersections between D and the Whitney disk interiors together with the framing obstructions on all the Whitney disks is equal to $\mathrm{Arf}(K)$ (see [**17**, **18**, **35**] and sketch just below). By performing boundary-twists on the Whitney disks as in Figure 19 (each of which changes a framing obstruction by ± 1), it can be arranged that all intersections between D and the Whitney disk interiors come in canceling pairs. This means that $\mathrm{Arf}(K)$ is now equal to the sum modulo 2 of the twistings on all the order 1 Whitney disks and that all order 1 intersections can be paired by order 2 Whitney disks. So K bounds an order 2 twisted Whitney tower $\mathcal{W}$ with $\mathrm{Arf}(K) = \tau_2^{\infty}(\mathcal{W})$ which counts the $(1,1)^{\infty}$ in $\mathcal{T}_2^{\infty}(1) \cong \mathbb{Z}_2$. On the other hand, given an arbitrary order 2 twisted Whitney tower $\mathcal{W}$ bounded by K, one has $\mathrm{Arf}(K) = \tau_2^{\infty}(\mathcal{W}) \in \mathcal{T}_2^{\infty}(1) \cong \mathbb{Z}_2$ determined again as the sum modulo 2 of twistings on all order 1 Whitney disks.

We sketch here a proof that $\mathrm{Arf}(K)$ is equal to the sum modulo 2 of the order 1 intersections plus framing obstructions in any *weak* order 1 Whitney tower $\mathcal{W} \subset B^4$ bounded by $K \subset S^3$. Here 'weak' means that the Whitney disks are not necessarily framed. (We are assuming that the Whitney disk boundaries are disjointly embedded, although we could instead also count Whitney disks boundary singularities.) Any K bounds a Seifert surface $F \subset S^3$, and by definition $\mathrm{Arf}(K)$ equals the sum modulo 2 of the products of twistings on dual pairs of 1-handles of F. Restricting to the case where F is genus 1, denote by γ and γ' core circles of the pair of dual 1-handles of F, with respective twistings a and a', so that $\mathrm{Arf}(K)$ is the product aa' modulo 2. Let D_γ be any immersed disk bounded by γ into B^4, so that the interior of D_γ is disjoint from F. After performing $|a|$ boundary-twists on D_γ, each of which creates a single intersection between D_γ and F, it can be arranged that D_γ is framed with respect to F, so that surgering F along D_γ creates only canceling pairs of self-intersections in the resulting disk D bounded by K. Each self-intersection in D_γ before the surgery contributes *two* canceling pairs of self-intersections of D, since the surgery adds both D_γ and an oppositely oriented parallel copy of D_γ to create D. On the other hand, the $|a|$ intersections between D_γ and F before the surgery give rise to exactly $|a|$ canceling pairs of self-intersections of D, so the total number of canceling pairs of self-intersections of D is equal to a modulo 2. Observe that all of these canceling pairs admit Whitney disks constructed from parallel copies of any immersed disk bounded by γ' with interior in B^4. The framing obstruction on each of these Whitney disks is equal to the twisting a' along γ', and the only order 1 intersections between the Whitney disk interiors and D come in canceling pairs, since they correspond to intersections with D_γ and its parallel copy. Thus, the sum of framing obstructions and order 1 intersections is equal to the product aa' modulo 2. The higher genus case is similar. That this construction is independent of the choice of weak Whitney tower follows from the fact that the analogous homotopy invariant for 2-spheres in 4-manifolds vanishes on any immersed 2-sphere in the 4-sphere (for example, [**46**], or [**18**, 10.8A and 10.8B]).

Considering now the second statement of Lemma 10 regarding links, it follows from Corollary 7 and Proposition 9 that if L is any link in $\mathrm{Ker}(\mu_2) < \mathsf{W}_2^{\infty}$, and $\mathcal{W}$ is any order 2 twisted Whitney tower bounded by L; then $\tau_2^{\infty}(\mathcal{W})$ is contained in the subgroup of $\mathcal{T}_2^{\infty}$

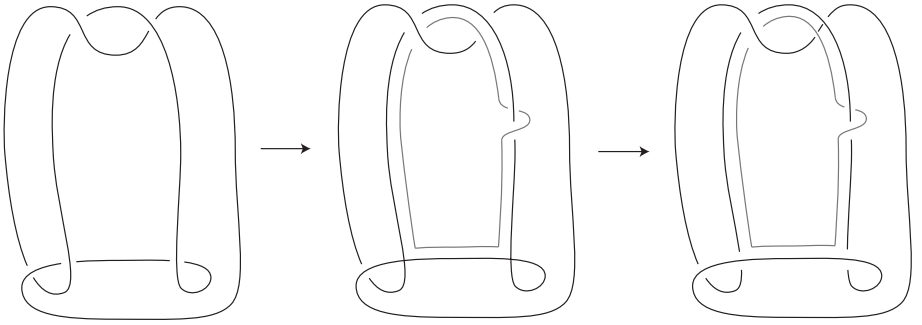

FIGURE 20. *From left to right*: *The trace of a null-homotopy of the figure-eight knot describes an order* 0 *disk* D *with a canceling pair of self-intersections that are paired by a clean* $+1$*-twisted Whitney disk* W*. A collar of* W *is indicated by the gray loop in the middle diagram, and the unlink in the right-hand diagram can be capped off by two embedded disks which form the rest of* D *and* W*. The twisting* $\omega(W) = 1$ *of* W *corresponds to the twist in the collar annulus in the middle diagram, as explained in Figure 21.*

spanned by the symmetric twisted trees $(i,i)^\infty$, and this subgroup is isomorphic to $(\mathbb{Z}_2)^m$. By the first statement of Lemma 10, the desired isomorphism $\mathrm{Arf}(L)$ is given by $\tau_2^\infty(\mathcal{W})$.

So to complete the proof of Lemma 10, it suffices to show that for any $L \in \mathrm{Ker}(\mu_2) < \mathsf{W}_2^\infty$, $L = 0 \in \mathsf{W}_2^\infty$ if and only if $\tau_2^\infty(\mathcal{W}) = 0$. But if $L = 0 \in \mathsf{W}_2^\infty$, then by definition L bounds an order 3 twisted Whitney tower, so $\tau_2^\infty(\mathcal{W}) = 0$. And if $\tau_2^\infty(\mathcal{W}) = 0$, then L bounds an order 3 twisted Whitney tower by Theorem 27, and hence $L = 0 \in \mathsf{W}_2^\infty$. □

6. *Boundary links and higher-order Arf invariants*

This section contains proofs of Lemma 13 and Proposition 14 from the introduction.

6.1. *Proof of Lemma* 13

Recall the statement of Lemma 13: For any rooted tree J of order $k-1$, by performing iterated untwisted Bing-doublings and interior band sums on the figure-eight knot K, one can create a boundary link K^J as the boundary of a twisted Whitney tower $\mathcal{W}$ of order $4k-2$ with $\tau_{4k-2}^\infty(\mathcal{W}) = \infty\!\!-\!\!\prec_J^J$.

First of all, Figures 20 and 21 show that the figure-eight knot K bounds an order 2 twisted Whitney tower consisting of an order 0 disk D which contains a single canceling pair of self-intersections, and a clean $+1$-twisted Whitney disk W pairing the self-intersections of D. (This exhibits the fact that K has non-trivial classical Arf invariant, as in Lemma 10; and the same computation shows that the Whitney disk in Figure 3 is $+1$-twisted.)

The link $\mathrm{Bing}(K)$ pictured in the left-hand side of Figure 22 is the untwisted Bing-double of the figure-eight knot, where 'untwisted' refers to the untwisted band with core K used to guide the construction of the two clasped unknotted components. The right-hand side of Figure 22 shows how the untwisted Bing-double of any boundary link is again a boundary link, as disjoint Seifert surfaces for the new pair of components can be constructed by banding together four parallel copies of the Seifert surface for the original component which was doubled.

Moving into B^4, a null-homotopy of $\mathrm{Bing}(K)$ which pulls apart the clasps (and is supported near the untwisted band) describes embedded order 0 disks D_1 and D_2 which have a single canceling pair of intersections (corresponding to the crossing-changes undoing the clasps). The boundary of an order 1 Whitney disk $W_{(1,2)}$ pairing $D_1 \cap D_2$ sits in an S^3-slice of B^4 as a figure-eight knot, exactly as in the left-hand side of Figure 20, with the rest of $W_{(1,2)}$

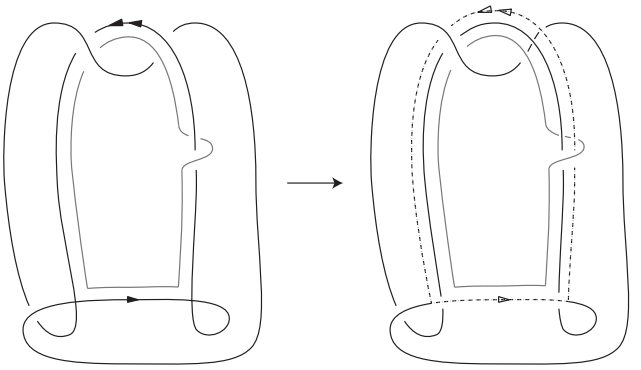

FIGURE 21. *The two right-most pictures from Figure 20 with the Whitney disk boundary arcs indicated by the arrows on the left and the corresponding arcs of the Whitney section shown by the dotted arcs on the right. (Compare with the framed Whitney disk in Figure 6.) The* $+1$*-linking between the gray circle (boundary of a collar of* W*) and the dotted circle (Whitney section) corresponds to the twisting* $\omega(W) = 1$.

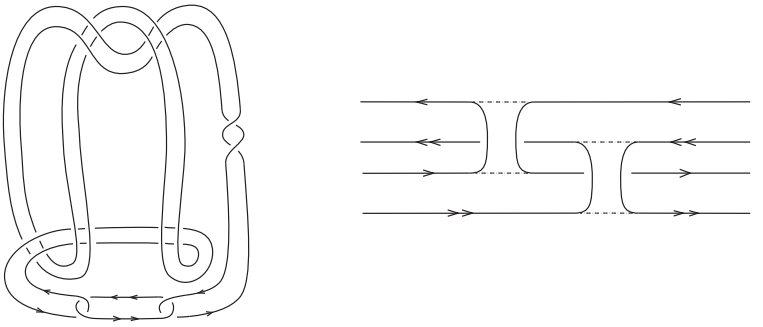

FIGURE 22. *On the left, the (untwisted) Bing-double of the figure-eight knot. On the right, disjoint Seifert surfaces for the two components of a Bing-doubled knot can be constructed from four copies of a Seifert surface bounded by the original knot.*

described by the same null-homotopy as for D of Figure 20. Since $\mathrm{Bing}(K)$ is untwisted, $W_{(1,2)}$ is framed [**10**, Section 3]. The interior of $W_{(1,2)}$ is disjoint from D_1 and D_2, but contains a canceling pair of self-intersections corresponding to the self-intersections of D in Figure 20. These self-intersections can be paired by an order 3 clean $+1$-twisted Whitney disk (the W in Figure 20) which has ∞-tree $((1,2),(1,2))^\infty$. Thus, $K^J := \mathrm{Bing}(K)$ for $J = (1,2)$ bounds an order 6 twisted Whitney tower $\mathcal{W}$ with $\tau_6^\infty(\mathcal{W}) = ((1,2),(1,2))^\infty \in \mathcal{T}_6^\infty$. As mentioned in the introduction, Conjecture 12 would imply that the $\mathrm{Bing}(K)$ does *not* bound any order 7 (twisted) Whitney tower $\mathcal{W}'$ on immersed disks in the 4-ball. Note that if $\mathrm{Bing}(K)$ did indeed bound such a $\mathcal{W}'$, then taking the union of $\mathcal{W}$ and $\mathcal{W}'$ along $\mathrm{Bing}(K)$ would yield a pair of immersed 2-spheres in the 4-sphere supporting the order 6 twisted Whitney tower $\mathcal{V}$ with $\tau_6(\mathcal{V}) = ((1,2),(1,2))^\infty$. This would imply that $\mathrm{Arf}_2(L) = 0$ for any link L since by tubing these 2-spheres as needed into any order 6 twisted Whitney tower bounded by L, one could kill any 2-torsion elements $((i,j),(i,j))^\infty \in \mathcal{T}_6^\infty$.

The construction for arbitrary J follows inductively by observing that having realized $(J,J)^\infty$ by K^J, with J of order r, Bing-doubling a component of K^J realizes $(J',J')^\infty$, with J' of order $r+1$ gotten from J by adding two new edges to a univalent vertex of J (see Figure 23); and any J' of order $r+1$ can be gotten from some such J, by banding together some link components as necessary to get repeated univalent labels. For instance, to realize $((1,(1,2)),(1,(1,2)))^\infty$ from

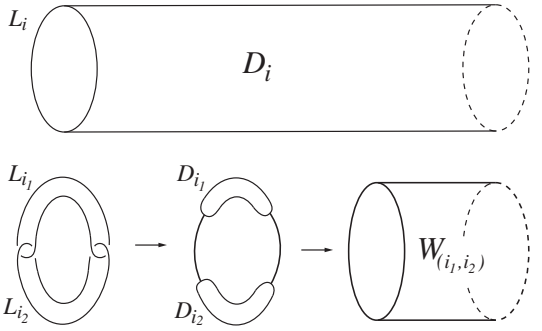

FIGURE 23. *Above, a collar of L_i in D_i (shown 'schematically' with L_i an unknot). Below (shown moving radially into B^4 from left to right), D_i can be converted into a Whitney disk $W_{(i_1,i_2)}$ pairing the intersections between order 0 disks D_{i_1} and D_{i_2} bounded by the components L_{i_1} and L_{i_2} of the untwisted Bing double of L_i. D_{i_1} and D_{i_2} are traced out by a null-homotopy of $L_{i_1} \cup L_{i_2}$ which 'pulls apart the clasp along L_i'.*

$((1,2),(1,2))^\infty$ realized by Bing(K) above, just Bing-double the second component of Bing(K) and then band one of the new components into the old first component.

6.2. *Proof of Proposition* 14

As stated in Conjecture 12, we believe that Arf_k is non-trivial for all k; however, interest in the first unknown 'test case' $k = 2$ is heightened by Proposition 14 from the introduction which states that if $\mathrm{Arf}_2 = 0$, then Arf_k is trivial for all $k \geqslant 2$.

By Proposition 9, it suffices to show that if the untwisted Bing-double $\mathrm{Bing}(K) = K^{((1,2),(1,2))}$ bounds an order 7 twisted Whitney tower for some K with non-trivial classical Arf invariant, then each link K^J of Lemma 13 with J of order $k-1$ bounds an order $4k-1$ twisted Whitney tower, for $k > 2$.

Note that the assumption that Bing(K) bounds an order 7 twisted Whitney tower implies that Bing(K) in fact bounds an order 10 twisted Whitney tower $\mathcal{W}$ by Theorem 8, since boundary links have vanishing Milnor invariants in all orders. Now applying the Bing-doubling and banding construction of the proof of Lemma 13 to get $K^{J'}$ from Bing(K), where J' is any order 2 tree gotten from the order 1 tree $J = (1,2)$, yields $K^{J'}$ bounding an order 11 twisted Whitney tower gotten from $\mathcal{W}$ by converting an order 0 disk of $\mathcal{W}$ into an order 1 Whitney disk (Figure 23). Inductively, if K^J, with J of order $k-1$, bounds an order $4k-1$ twisted Whitney tower, then K^J also bounds an order $4k+2$ twisted Whitney tower by Theorem 8, and since Bing-doubling a component of K^J raises the order by at least 1, it follows that $K^{J'}$ bounds an order $4(k+1)-1$ twisted Whitney tower, for any J' gotten from J by attaching at least one pair of new edges to a univalent vertex of J. As was observed in the proof of Lemma 13, all trees can be gotten by this process of adding new pairs of edges to univalent vertices of lower-order trees. □

7. *Milnor invariants and geometric k-sliceness*

This section gives proofs of Theorems 16 and 18 from the introduction. The proof of Theorem 16 uses the classification of the twisted Whitney tower filtration from [**10**], together with the Whitney tower-to-grope techniques of [**44**] (as sketched in Section 3). The proof of Theorem 18 will use Theorem 17 of the introduction, together with a mild generalization of Theorem 6 given by Proposition 34.

7.1. *Proof of Theorem 16*

Recall from Subsection 1.6 that Theorem 16 states: a link L is geometrically k-slice if and only if $\mu_n(L) = 0$ for all $n \leqslant 2k-2$ and $\mathrm{Arf}_n(L) = 0$ for all $n \leqslant k$; where $L \subset S^3$ is *geometrically k-slice* if the components L_i bound disjointly embedded (oriented) surfaces $\Sigma_i \subset B^4$ such that a symplectic basis of curves on each Σ_i bound disjointly embedded framed class k gropes in the complement of $\Sigma := \bigcup_i \Sigma_i$.

Proof. By the classification of the twisted Whitney tower filtration [**10**, Corollary 1.16], the stated vanishing of $\mu_n(L)$ and $\mathrm{Arf}_n(L)$ is equivalent to L bounding an order $2k-1$ twisted Whitney tower in B^4 (see Theorem 8, Proposition 9 and Definition 11). So it will suffice to show that L is in $\mathbb{W}^{\infty}_{2k-1}$ if and only if L is geometrically k-slice.

Recall that $\mathbb{W}^{\infty}_{2k-1} = \mathbb{W}_{2k-1}$ by definition, so we may assume that L bounds a framed Whitney tower $\mathcal{W}$ of order $2k-1$. By applications of the Whitney-move IHX construction [**45**, Section 7] it can be arranged that all trees in the intersection forest $t(\mathcal{W})$ are *simple*, meaning that every trivalent vertex is adjacent to at least one univalent edge. Since all these simple trees are of order (at least) $2k-1$, we can choose a preferred univalent vertex on each tree which is (at least) $k-1$ trivalent vertices away from both ends of its tree. Now converting the order $2k-1$ Whitney tower $\mathcal{W}$ to a class $2k$ embedded grope G via (the framed part of) the above construction in the proof of Lemma 31 (as described in detail in [**45**]), with the preferred univalent vertices corresponding to the bottom stages of the connected components of G, yields dyadic branches having bottom stages with a symplectic basis of circles bounding gropes of class (at least) k (each tree yields one hyperbolic pair of such circles).

Note that the construction of [**45**] used here, as in the proof of Lemma 31, yields a *capped* grope G^c which is contained in any small neighborhood of $\mathcal{W}$. In this argument, we only need the body G.

On the other hand, being geometrically k-slice is the same as bounding a particular kind of embedded class $2k$ grope $G \subset B^4$. Since B^4 is simply connected, caps can be found, and can be framed by twisting as necessary. All intersections in the caps can be pushed down into the bottom grope stages using finger moves, yielding a capped grope G^c bounded by the link, which can be converted into an order $2k-1$ Whitney tower via the inverse operation to that used in the proof of Lemma 31 (see [**45**, Theorem 6]). □

7.2. *Proof of Theorem* 18

Recall the statement of Theorem 18: A link $L = \bigcup_i L_i$ has $\mu_n(L) = 0$ for all $n \leqslant 2k-2$ if and only if the link components L_i 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 a symplectic basis of curves on Σ'_i bound disjointly embedded framed gropes $G_{i,j}$ of class k in the complement of $\Sigma := \bigcup_i \Sigma_i$, and 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 $G_{i,j}$.

Proof. Given L with vanishing Milnor invariants of all orders $\leqslant 2k-2$, by Theorem 17 there exist finitely many boundary links as in Lemma 13 such that taking band sums of L with all these boundary links yields a geometrically k-slice link $L' \subset S^3$. Consider each of these boundary links to be contained in a 3-ball, and embed these 3-balls disjointly in a single 3-sphere, so the union of the boundary links forms a single boundary link denoted by U. Decompose the 3-sphere $S^3 = B^3_L \cup_{S^2_0} B^3_U$ containing L' into two 3-balls exhibiting the band sum of $L' = L \# U$, with $L \subset B^3_L$ and $U \subset B^3_U$. Each band in the sum intersects the separating 2-sphere S^2_0 in a single transverse arc. Since L' is geometrically k-slice, L' bounds $\Sigma' \subset B^4$, which satisfies the conditions in the first item of Theorem 18.

Now consider $S^3 = B^3_L \cup_{S^2_0} B^3_U$ as the equator of a 4-sphere S^4, with the interior of $\Sigma' \subset B^4$ contained in the 'southern hemisphere' $B^4 \subset S^4$. The components of the boundary link U bound disjoint Seifert surfaces which are contained in B^3_U, and symplectic bases of these Seifert surfaces bound immersed disks into the 'northern hemisphere' 4-ball in S^4. We may assume that the interiors of these immersed disks are contained in a 'northern quadrant' of S^4, which is a 4-ball B^4_+ bounded by a 3-sphere consisting of $B^3_+ \cup_{S^2_0} B^3_U$, where B^3_+ is a 3-ball bounded by S^2_0 whose interior cuts the northern hemisphere into two 4-balls. Gluing B^4_+ to the southern hemisphere B^4 along B^3_U, with the boundaries of the Seifert surfaces glued along U, has the effect of eliminating U from the band sum with L (U gets replaced by the unlink). This leaves L in the 3-sphere $B^3_+ \cup_{S^2_0} B^3_+$ which bounds the 'other' northern quadrant of S^4, and the union Σ'' of the Seifert surfaces together with Σ' form the surfaces Σ as desired.

Conversely, suppose the components L_i of L bound disjointly embedded surfaces $\Sigma_i \subset B^4$ as in the statement. The class k gropes $G_{i,j}$ attached to dual circles in Σ'_i can be thought of as grope branches of class $2k$ by subdividing each Σ'_i into genus 1 pieces. Caps can be chosen for all tips of these branches, and by pushing down intersections (using finger moves), it can be arranged that the caps only intersect the bottom stages Σ'_i. This means that the caps are disjointly embedded, and disjoint from the immersed disks bounded by the symplectic bases in Σ''. Now applying the capped grope-to-Whitney tower construction of [**45**, Theorem 6] to these capped branches yields an order $2k-1$ Whitney tower $\mathcal{W}$ on immersed surfaces S_i each bounded by L_i such that all Whitney disks and singularities of $\mathcal{W}$ are contained in a neighborhood of the capped branches, and with $\Sigma''_i \subset S_i$ for each i. In particular, a symplectic basis on each S_i bounds immersed disks whose interiors are contained in the complement of $\mathcal{W}$.

The proof is completed by the following proposition, which mildly generalizes Theorem 6 and in particular implies that L as above has vanishing Milnor invariants of all orders at most $2k-2$. □

PROPOSITION 34. *Theorem* 6 *holds for an order* n *twisted Whitney tower* $\mathcal{W} \subset B^4$ *on order zero immersed surfaces* S_i *bounded by* L *such that a symplectic basis of curves on each* S_i *bounds immersed disks in the complement of* $\mathcal{W}$*: The Milnor invariants* $\mu_k(L)$ *vanish for* $k < n$*, and* $\mu_n(L) = \eta_n \circ \tau^{\infty}_n(\mathcal{W})$.

Proof. We work through each step of the proof of Theorem 6 given in Section 4, checking that the assertions still hold when the order 0 disks bounded by the link components are replaced by surfaces S_i.

The twisted Whitney tower $\mathcal{W}$ is resolved to a twisted capped grope G^c just as in the proof of Lemma 31 except that the bases of curves from the S_i are left uncapped. Note that G^c does not really have class $n+1$ because no higher grope stages are attached to these basis curves; however, we will see that the proof still goes through since these curves bound immersed disks in $B^4 \setminus G^c$.

To see that Lemma 33 still holds, the only new point that needs to be checked in the proof given in Subsection 4.4 is that the new generators of $H_2(B^4 \setminus G^c)$ which are Alexander dual to the basis curves on the S_i are represented by maps of gropes of class at least $n+2$. These new generators are in fact represented by maps of 2-spheres (which are gropes of arbitrary high class): A torus consisting of the union of circle fibers in the normal circle bundle over a basis curve contains a dual pair of circles, one of which is a meridian to S_i (and bounds a normal disk to the basis curve, exhibiting Alexander duality), while the other circle (which is parallel to the basis curve) bounds by assumption an immersed disk in the complement of G^c. Therefore, each such torus can be surgered to an immersed 2-sphere in $B^4 \setminus G^c$.

It only remains to check that the computation of the link longitudes in Subsection 4.2 still corresponds to the composition $\eta_n(\tau^{\infty}_n(\mathcal{W}))$. But this is clear since all the basis curves from the S_i represent trivial elements in $\pi_1(B^4 \setminus G^c)$. □

Acknowledgements. 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. Thanks also to the referee for helpful comments on the exposition.

References

1. D. Bar-Natan, 'Vassiliev homotopy string link invariants', *J. Knot Theory Ramifications* 4 (1995) 13–32.
2. A. Casson, 'Link cobordism and Milnor's invariant', *Bull. London Math. Soc.* 7 (1975) 39–40.
3. A, Casson, 'Three lectures on new constructions in 4-dimensional manifolds', *A la recherche de la topologie perdue*, Progress in Mathematics 62 (Birkhauser Boston, Boston MA, 1986) 201–244.
4. J. C. Cha, 'Link concordance, homology cobordism, and Hirzebruch-type defects from iterated p-covers', *J. Eur. Math. Soc.* 12 (2010) 555–610.
5. T. Cochran, 'Derivatives of links, Milnor's concordance invariants and Massey products', *Mem. Amer. Math. Soc.* 84 (1990).
6. T. Cochran, 'k-cobordism for links in S^3', *Trans. Amer. Math. Soc.* 327 (1991) 641–654.
7. T. Cochran and P. Melvin, 'The Milnor degree of a 3-manifold', *J. Topol.* 3 (2010) 405–423.
8. J. Conant, R. Schneiderman and P. Teichner, 'Jacobi identities in low-dimensional topology', *Compos. Math.* 143 Part 3 (2007) 780–810.
9. J. Conant, R. Schneiderman and P. Teichner, 'Higher-order intersections in low-dimensional topology', *Proc. Natl Acad. Sci. USA* 108 (2011) 8131–8138.
10. J. Conant, R. Schneiderman and P. Teichner, 'Whitney tower concordance of classical links', *Geom. Topol.* 16 (2012) 1419–1479.
11. J. Conant, R. Schneiderman and P. Teichner, 'Tree homology and a conjecture of Levine', *Geom. Topol.* 16 (2012) 555–600.
12. J. Conant and P. Teichner, 'Grope coborbism of classical knots', *Topology* 43 (2004) 119–156.
13. J. Conant and P. Teichner, 'Grope cobordism and Feynman diagrams', *Math. Ann.* 328 (2004) 135–171.
14. S. Donaldson, 'An application of gauge theory to four-dimensional topology', *J. Differential Geom.* 18 (1983) 279–315.
15. W. Dwyer, 'Homology, Massey products and maps between groups', *J. Pure Appl. Algebra* 6 (1975) 177–190.
16. M. Freedman, 'The topology of four-dimensional manifolds', *J. Differential Geom.* 17 (1982) 357–453.
17. M. Freedman and R. Kirby, 'A geometric proof of Rochlin's theorem', *Algebraic and geometric topology (Proc. Sympos. Pure Math., Stanford Univ., Stanford, Calif., 1976), Part 2*, Proceedings of Symposia in Pure Mathematics XXXII (American Mathematical Society, Providence, RI, 1978) 85–97.
18. M. Freedman and F. Quinn, *The topology of 4-manifolds*, Princeton Mathematical Series 39 (Princeton University Press, Princeton, NJ, 1990).
19. M. Freedman and P. Teichner, '4-manifold topology II: Dwyer's filtration and surgery kernels', *Invent. Math.* 122 (1995) 531–557.
20. S. Gadgil, 'Watson–Crick pairing, the Heisenberg group and Milnor invariants', *J. Math. Biol.* 59 (2009) 123–142.
21. N. Habegger and X. Lin, 'The classification of links up to homotopy', *J. Amer. Math. Soc.* 3 (1990) 389–419.
22. N. Habegger and X. Lin, 'On link concordance and Milnor's μ-invariants', *Bull. London Math. Soc.* 30 (1998) 419–428.
23. N. Habegger and G. Masbaum, 'The Kontsevich integral and Milnor's invariants', *Topology* 39 (2000) 1253–1289.
24. N. Habegger and K. Orr, 'Milnor link invariants and quantum 3-manifold invariants', *Comment. Math. Helv.* 74 (1999) 322–344.
25. K. Habiro and J.-B. Meilhan, 'Finite type invariants and Milnor invariants for Brunnian links', *Int. J. Math.* 19 (2008) 747–766.
26. J. Hillman, *Algebraic invariants of links*, Series on Knots and everything 32 (World Scientific, Singapore, 2002).
27. K. Igusa and K. Orr, 'Links, pictures and the homology of nilpotent groups', *Topology* 40 (2001) 1125–1166.
28. M. Kervaire and J. Milnor, ' On 2-spheres in 4-manifolds', *Proc. Natl Acad. Sci.* 47 (1961) 1651–1657.
29. V. Krushkal, 'Exponential separation in 4-manifolds', *Geom. Topol.* 4 (2000) 397–405.
30. V. Krushkal and P. Teichner, 'Alexander duality, gropes and link homotopy', *Geom. Topol.* 1 (1997) 51–69.
31. J. Levine, 'Homology cylinders: an enlargement of the mapping class group', *Algebr. Geom. Topol.* 1 (2001) 243–270.
32. J. Levine, 'Addendum and correction to: Homology cylinders: an enlargement of the mapping class group', *Algebr. Geom. Topol.* 2 (2002) 1197–1204.
33. J. Levine, 'Labeled binary planar trees and quasi-Lie algebras', *Algebr. Geom. Topol.* 6 (2006) 935–948.
34. W. Magnus, A. Karass and D. Solitar, *Combinatorial group theory* (Dover Publications, Inc., New York, 1976).

35. Y. Matsumoto, 'Secondary intersectional properties of 4-manifolds and Whitney's trick', *Proc. Sympos. Pure Math.* 32, Part 2 (1978) 99–107.
36. J.-B. Meilhan and A. Yasuhara, 'Milnor invariants and the HOMFLYPT polynomial', *Geom. Topol.* 16 (2012) 889–917.
37. J. Milnor, 'Link groups', *Ann. of Math.* 59 (1954) 177–195.
38. J. Milnor, 'Isotopy of links', *Algebraic geometry and topology* (Princeton University Press, Princeton, NJ, 1957).
39. M. Morishita, 'Milnor invariants and Massey products for prime numbers', *Compos. Math.* 140, Part 1 (2004) 69–83.
40. K. Orr, 'Homotopy invariants of links', *Invent. Math.* 95 (1989) 379–394.
41. R. Porter, 'Milnor's $\overline{\mu}$-invariants and Massey products', *Trans. Amer. Math. Soc.* 257 (1980) 39–71.
42. V. Rohlin, 'New results in the theory of four-dimensional manifolds', *Dokl. Akad. Nauk SSSR (N.S.)* 84 (1952) 221–224 (Russian).
43. N. Sato, 'Cobordisms of semi-boundary links', *Topology Appl.* 18 (1984) 225–234.
44. R. Schneiderman, 'Simple Whitney towers, half-gropes and the Arf invariant of a knot', *Pacific J. Math.* 222 (2005) 169–184.
45. R. Schneiderman, 'Whitney towers and gropes in 4-manifolds', *Trans. Amer. Math. Soc.* 358 (2006) 4251–4278.
46. R. Schneiderman and P. Teichner, 'Higher order intersection numbers of 2-spheres in 4-manifolds', *Algebr. Geom. Topol.* 1 (2001) 1–29.
47. R. Schneiderman and 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 7 (Geometry and Topology Publications, Coventry, 2004) 101–134.
48. A. Scorpan, *The wild world of 4-manifolds* (American Mathematical Society, Providence, RI, 2005).
49. J. Stallings, 'Homology and central series of groups', *J. Algebra* 2 (1965) 170–181.
50. P. Teichner, 'Knots, von Neumann signatures, and grope cobordism', *Proceedings of the International Congress of Math. Vol II: Invited Lectures* (ICM 2002) 437–446.
51. P. Teichner, 'What is... a grope?', *Notices Amer. Math. Soc.* 54 (2004) 894–895.
52. V. G. Turaev, 'Milnor invariants and Massey products', *Investigations in Topology. Part II, Zap. Nauchn. Sem. LOMI 66, "Nauka"* (Leningrad. Otdel., Leningrad, 1976) 189–203.
53. H. Whitney, 'The self intersections of a smooth n-manifold in $2n$-space', *Ann. of Math.* 45 (1944) 220–246.

J. Conant
Department of Mathematics
University of Tennessee
Knoxville, TN 37996-1320
USA

jconant@math.utk.edu

R. Schneiderman
Department of Mathematics and Computer Science
Lehman College
City University of New York
Bronx, NY 10468
USA

robert.schneiderman@lehman.cuny.edu

P. Teichner
Department of Mathematics
University of California
Berkeley, CA 94720-3840
USA

and

Max-Planck-Institut für Mathematik
Bonn
Germany

teichner@mac.com