Higher-Order Intersections in Low-Dimensional Topology

Jim Conant [*], **Rob Schneiderman** [†], **and Peter Teichner** [‡]

[*]University of Tennessee,[†]Lehman College CUNY, and [‡]UC Berkeley and Max-Planck-Institut für Mathematik, Bonn

Submitted to Proceedings of the National Academy of Sciences of the United States of America

We show how to measure the failure of the Whitney move in dimension 4 by constructing higher-order intersection invariants of *Whitney towers* built from iterated Whitney disks on immersed surfaces in 4–manifolds. For Whitney towers on immersed disks in the 4–ball, we identify some of these new invariants with previously known link invariants like Milnor, Sato-Levine and Arf invariants. We also define higher-order Sato-Levine and Arf invariants and show that these invariants detect the obstructions to framing a *twisted* Whitney tower. Together with Milnor invariants, these higher-order invariants are shown to classify the existence of (twisted) Whitney towers of increasing order in the 4–ball. A conjecture regarding the non-triviality of the higher-order Arf invariants is formulated, and related implications for filtrations of string links and 3–dimensional homology cylinders are described.

Whitney tower | Higher-order intersection | 4–manifold | link concordance | trivalent tree | Arf invariant | k-slice | Gropes

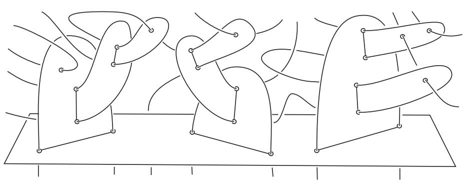

Fig. 1. Part of a Whitney tower in 4–space.

Introduction

Despite how it may appear in high school, mathematics is not all about manipulating numbers or functions in more and more complicated algebraic or analytic ways. In fact, one of the most interesting quests in mathematics is to find a good notion of *space*. It should be general enough to cover many real life situations and at the same time sufficiently specialized so that one can still prove interesting properties about it. A first candidate was Euclidean n-space $\mathbb{R}^n$, consisting of n-tuples of real numbers. This covers all dimensions n but is too special: the surface of the earth, mathematically modelled by the 2–sphere S^2, is 2-dimensional but compact, so it can't be $\mathbb{R}^2$. However, S^2 is *locally* Euclidean: around every point one can find a neighborhood which can be completely described by two real coordinates (but *global* coordinates don't exist).

This observation was made into the definition of an *n-dimensional manifold* in 1926 by Kneser: It's a (second countable) Hausdorff space which looks locally like $\mathbb{R}^n$. The development of this definition started at least with Riemann in 1854 and important contributions were made by Poincaré and Hausdorff at the turn of the 19th century. It covers many important physical notions, like the surface of the earth, the universe, and space-time (for $n = 2, 3,$ and 4, respectively) but is special enough to allow interesting structure theorems. One such statement is *Whitney's (strong) embedding theorem*: Any n-manifold M^n can be embedded into $\mathbb{R}^{2n}$ (for all $n \geq 1$). The proof in small dimensions $n = 1, 2$ is fairly elementary and special, but in all dimensions $n > 2$, Hassler Whitney [1] found the following beautiful argument: By general position, one finds an immersion $M \to \mathbb{R}^{2n}$ with at worst transverse double points. By adding local cusps, one can assume that all double points can be paired up by *Whitney disks* as in Figure 2, using the fact that $\mathbb{R}^{2n}$ is simply connected. Since $2+2 < 2n$ and $n+2 < 2n$, one can arrange that all Whitney disks are disjointly embedded, framed and meet the image of M only on the boundary. Then a sequence of *Whitney moves*, as shown in Figure 2, leads to the desired embedding of M.

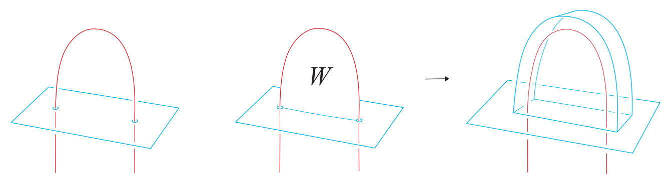

Fig. 2. Left: A canceling pair of transverse intersections between two local sheets of surfaces in a 3-dimensional slice of 4-space. The horizontal sheet appears entirely in the 'present', and the red sheet appears as an arc which is assumed to extend into 'past' and 'future'. Middle: A Whitney disk W pairing the intersections. Right: A Whitney move guided by W eliminates the intersections.

The Whitney move, sometimes also called the *Whitney trick*, remains a primary tool for turning algebraic information (counting double points) into geometric information (existence of embeddings). It was successfully used in the classification of manifolds of dimension > 4, specifically in Smale's celebrated h-cobordism theorem [2] (implying the Poincaré conjecture) and Wall's surgery theory [3]. The failure of the Whitney move in dimension 4 is the main reason that, even today, there is no classification of 4-dimensional manifolds in sight. To be more precise, one needs to distinguish between *topological* and *smooth* 4–manifolds to make correct statements. A topological n–manifold is locally *homeomorphic* to $\mathbb{R}^n$, whereas a smooth manifold is locally *diffeomorphic* to it (in the given smooth structure).

Casson realized that in the setting of the 4-dimensional h-cobordism theorem, even though Whitney disks can't always be embedded (because $2+2 = 4$), they always fit into what is now called a *Casson tower*. This is an iterated construction that works in simply connected 4–manifolds, where one adds more and more layers of disks onto the singularities of a given

Reserved for Publication Footnotes

(immersed) Whitney disk [4]. In an amazing tour de force, Freedman [5, 6] showed that there is always a *topologically* embedded disk in a neighborhood of certain Casson towers (originally, one needed 7 layers of disks, later this was reduced to 3). This result implied the topological h-cobordism theorem (and hence the topological Poincaré conjecture) in dimension 4. At the same time, Donaldson used gauge theory to show that the smooth 4-dimensional h-cobordism theorem fails [7], and both results were awarded with a Fields medal in 1982. Surprisingly, the smooth Poincaré conjecture is still open in dimension 4 – the only remaining unresolved case.

In the non-simply connected case, even the topological classification of 4–manifolds is far from being understood because Casson towers cannot always be constructed. See [8, 9, 10] for a precise formulation of the problem and a solution for fundamental groups of subexponential growth. However, there is a simpler construction, called a *Whitney tower*, which can be performed in many more instances. Here one again adds more and more layers of disks to a given (immersed) Whitney disk, however, one doesn't control all intersections as in a Casson tower but only *pairs* of intersections that allow higher order Whitney disks, see Figure 3 below. Thus a Casson tower gives a Whitney tower but not vice versa.

The current authors have developed an obstruction theory for such Whitney towers in a sequence of papers [11, 12, 13, 14, 15, 16, 17, 18, 19]. Even though the existence of a Whitney tower does not lead to an embedded (topological) disk, it is still a necessary condition. Hence our obstruction theory provides *higher-order (intersection) invariants* for the existence of embedded disks, spheres, or surfaces in 4–manifolds.

The easiest example of our intersection invariant is Wall's self-intersection number for disks in 4-manifolds. If $A : (D^2, \partial D^2) \to (M^4, \partial M)$ has trivial self-intersection number (we say that the *order zero invariant* $\tau_0(A)$ vanishes) then all self-intersections can be paired up by Whitney disks W_i. However, the W_i will in general self-intersect and intersect each other and also the original disk A. Our (first order) intersection invariant $\tau_1(A, W_i)$ counts the transverse intersections $A \pitchfork W_i$ and vanishes if they all can be paired up by (second order) Whitney disks $W_{i,j}$. This procedure continues with an invariant $\tau_2(A, W_i, W_{i,j})$ which measures both $A \pitchfork W_{i,j}$ and $W_i \pitchfork W_k$ intersections, and the construction of a higher-order Whitney tower $\mathcal{W}$ if the invariant vanishes. $\mathcal{W}$ is the union of A (at order 0) and all Whitney disks W_i (order 1), $W_{i,j}$ (order 2) and continuing with higher-order Whitney disks. If A is homotopic (rel. boundary) to an embedding then these constructions can be continued ad infinitum.

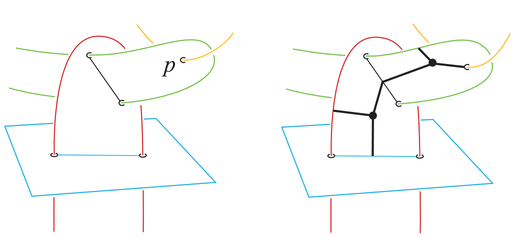

Fig. 3. An unpaired intersection point p among local sheets in a Whitney tower (left), and its associated tree (right).

The intersection invariants $\tau_n(A, W_i, W_{i,j}, \dots) = \tau_n(\mathcal{W})$ take values in a finitely generated abelian group $\mathcal{T}_n$ which is generated by certain trivalent trees that describe the 1-skeleton of a Whitney tower (Figure 3 and Definition 4). The relations in $\mathcal{T}_n$ correspond to Whitney moves, and quite surprisingly most of these relations can be expressed in terms of the so called IHX-relation which is a geometric incarnation of the Jacobi identity for Lie algebras. All the relations can be realized by controlled manipulations of Whitney towers, and as a result we recover the following approximation of the "algebra implies geometry" principle that is available in high dimensions:

Theorem 1. (Raising the order of a Whitney tower) *If A supports an order n Whitney tower $\mathcal{W}$ with vanishing $\tau_n(\mathcal{W})$, then A is homotopic (rel. boundary) to A' which supports an order $n+1$ Whitney tower. Compare Theorem 18.*

As usual in an obstruction theory, the dependence on the lower order Whitney towers makes it hard to derive explicit invariants that prevent the original disk A from being homotopic to an embedding. In this paper we discuss how to solve this problem in the easiest possible ambient manifold $M = B^4$, the 4-dimensional ball. We start with maps

$$A_1, \dots, A_m : (D^2, S^1) \to (B^4, S^3)$$

which exhibit a fixed link in the boundary 3–sphere S^3. If this link was *slice*, then the A_i would be homotopic (rel. boundary) to disjoint embeddings; and our Whitney tower theory gives obstructions to this situation. In the simplest example discussed above we have $m = 1$ and the boundary of A is just a knot K in S^3:

Theorem 2. (The easiest case of knots [13]) *The first order intersection invariant $\tau_1(A, W_i) \in \mathcal{T}_1 \cong \mathbb{Z}_2$ can be identified with the Arf invariant of the knot K. It is thus a well-defined invariant that only depends on $\partial A = K$. Moreover, it is the complete obstruction to finding a Whitney tower of arbitrarily high order ≥ 2 with boundary K.*

There is a very interesting refinement of the theory for knots in the setting of Cochran, Orr, and Teichner's *n-solvable filtration*: Certain special symmetric Whitney towers of orders which are powers of 2 have a refined measure of complexity called *height*, and are obstructed by higher-order signatures of associated covering spaces [20]. However there are no known algebraic criteria for "raising the height" of a Whitney tower analogous to Theorem 1.

If $m > 1$ then the order zero invariant $\tau_0(A_1, \dots, A_m)$ is given by the linking numbers of the components $L_i := \partial A_i$ of the link $L = \cup_{i=1}^m L_i \subset S^3$ that is the boundary of the given disks. Milnor [21, 22] showed in 1954 how to generalize linking numbers $\mu(i,j)$ inductively to higher order. Here we use the *order n total Milnor invariants μ_n* which correspond to all length $(n+2)$ Milnor numbers $\mu(i_1, \dots, i_{n+2})$.

Theorem 3. (Milnor numbers as intersection invariants) *If a link L bounds a Whitney tower $\mathcal{W}$ of order n then the Milnor invariants μ_k of order $k < n$ vanish. Moreover, the order n Milnor invariants of L can be computed from the* intersection invariant *$\tau_n(\mathcal{W}) \in \mathcal{T}_n$. Compare Theorem 20.*

In the remaining sections, we will make these statements precise and explain how to get *complete* obstructions for the existence of Whitney towers for links. Unlike the case of knots, these get more and more interesting for increasing order. In addition to the above Milnor invariants (higher-order linking numbers), we'll need higher-order versions of Sato-Levine and Arf invariants. In a fixed order, these are finitely many $\mathbb{Z}_2$-valued invariants, so that, surprisingly, the Milnor invariants already detect the problem up to this 2-torsion information.

Theorem 4. (Classification of Whitney tower concordance) *A link L bounds a Whitney tower $\mathcal{W}$ of order n if and only if its Milnor invariants, Sato-Levine invariants and Arf invariants vanish up to order n. Compare Corollary 10.*

To prove this classification, we use Theorem 1 to show that the intersection invariant $\tau_n(\mathcal{W})$ leads to a surjective *realization map* $R_n : \mathcal{T}_n \twoheadrightarrow \mathsf{W}_n$, where W_n consist of links bounding

Whitney towers of order n, up to order $n+1$ Whitney tower concordance (see the next section). The Milnor invariant can be translated into a homomorphism $\mu_n : \mathsf{W}_n \to \mathsf{D}_n$, where the latter is a group defined from a free Lie algebra (which can be expressed via *rooted* trivalent trees modulo the Jacobi identity). The composition

$$\eta_n : \mathcal{T}_n \to \mathsf{W}_n \to \mathsf{D}_n$$

is hence a map between purely combinatorial objects both given in terms of trivalent trees. Using a geometric argument (grope duality), we show that it is simply given by summing over all choices of a root in a given tree (which is a more precise statement of Theorem 3). This map was previously studied by Jerry Levine in his work on 3-dimensional homology cylinders [23, 24], where he made a precise conjecture about the kernel and cokernel of η_n. He verified the conjecture for the cokernel in [25], using a generalized Hall algorithm.

In [14] we prove Levine's full conjecture via an application of combinatorial Morse theory to tree homology. In particular, we show that the kernel of η_n consists only of 2-torsion. This 2-torsion corresponds to our higher-order Sato-Levine and Arf invariants, and is characterized geometrically in terms of a framing obstruction for *twisted* Whitney towers (in which certain Whitney disks are not required to be framed).

In the above classification of Whitney tower concordance there remains one key geometric question: Although our higher-order Arf invariants are well-defined, it is not currently known if they are in fact non-trivial. All potential values are realized by simple links, so the question here is whether or not there are any further geometric relations; see Definition 2. We conjecture that indeed all the higher-order Arf invariants are non-trivial, or equivalently, that our realization maps $\widetilde{R}_n : \widetilde{\mathcal{T}}_n \to \mathsf{W}_n$ are isomorphisms for all n. Here $\widetilde{\mathcal{T}}_n$ is a certain quotient of $\mathcal{T}_n$ by what we call *framing relations* which come from IHX-relations on twisted Whitney towers. For $n \equiv 0, 2, 3 \mod 4$ we do show that $\widetilde{R}_n$ is an isomorphism, implying that in this further quotient the intersection invariant $\tau_n(\mathcal{W})$ only depends on the link $\partial\mathcal{W}$, and not on the choice of Whitney tower $\mathcal{W}$. The higher-order Arf invariants appear when $n = 4k-3$, and our conjecture says that the same conclusion holds in these orders.

This conjecture is in turn equivalent to the vanishing of the intersection invariants on all immersed 2–spheres in S^4. Of course all such maps are null-homotopic, and a general goal of the Whitney tower theory is to extract higher-order invariants of representatives of classes in the second homotopy group $\pi_2 M$. This obstruction theory is still being developed but certain aspects of it appeared in [11, 18, 19, 26]. The fundamental group $\pi_1 M$ leads to more interesting obstruction groups $\mathcal{T}_n(\pi_1 M)$ and a non-trivial $\pi_2 M$ leads to more relations to make the intersection invariants only dependent on the order zero surfaces.

In this paper, we will give a survey of the material needed to understand the above results for Whitney towers in the 4–ball. More details and proofs can be found in our recent series of five papers [12, 13, 14, 15, 16] from which we'll also survey here the following aspects of the theory:

- Twisted Whitney towers and their framing obstructions
- Geometrically k-slice links and vanishing Milnor invariants
- String links and the Artin representation
- Filtrations of 3-dimensional homology cylinders

Whitney towers

We work in the smooth oriented category (with discussions of orientations mostly suppressed), even though all results hold in the locally flat topological category by the basic results on topological immersions in Freedman–Quinn [8]. In particular, as remarked in [12] our techniques do not distinguish smooth from locally flat surfaces.

Order n Whitney towers are defined recursively as follows.

Definition 1. *A* surface of order 0 *in an oriented 4–manifold M is a connected oriented surface in M with boundary embedded in the boundary and interior immersed in the interior of M. A* Whitney tower of order 0 *is a collection of order 0 surfaces. The* order of a (transverse) intersection point *between a surface of order n and a surface of order m is $n+m$. The* order of a Whitney disk *is $(n+1)$ if it pairs intersection points of order n. For $n \geq 1$, a* Whitney tower of order n *is a Whitney tower $\mathcal{W}$ of order $(n-1)$ together with (immersed) Whitney disks pairing all order $(n-1)$ intersection points of $\mathcal{W}$.*

The Whitney disks in a Whitney tower may self-intersect and intersect each other as well as lower order surfaces but the boundaries of all Whitney disks are required to be disjointly embedded. In addition, all Whitney disks are required to be *framed*, as will be discussed below.

Whitney tower concordance. We now specialize to the case $M = B^4$, and also assume that a Whitney tower $\mathcal{W}$ has disks for its order 0 surfaces which have an m-component link in $S^3 = \partial B^4$ as their boundary, denoted $\partial\mathcal{W}$. Let $\mathbb{W}_n$ be the set of all framed links $\partial\mathcal{W}$ where $\mathcal{W}$ is an order n Whitney tower, and the link framing is induced by the order 0 disks in $\mathcal{W}$. This defines a filtration $\cdots \subseteq \mathbb{W}_3 \subseteq \mathbb{W}_2 \subseteq \mathbb{W}_1 \subseteq \mathbb{W}_0 \subseteq \mathbb{L}$ of the set of framed m-component links $\mathbb{L} = \mathbb{L}(m)$. Note that $\mathbb{W}_0$ consists of links that are evenly framed because a component has even framing if and only if it bounds a framed immersed disk in B^4.

In order to detect what stage of the filtration a particular link lies in, it would be convenient to define a set measuring the difference between $\mathbb{W}_n$ and $\mathbb{W}_{n+1}$. Because these are sets and not groups, the quotient is not defined. However we can still define an associated graded set in the following way.

Suppose $\mathcal{W}$ is an order $n+1$ Whitney tower in $M = S^3 \times [0,1]$ where each of the order 0 surfaces $A_1, \ldots, A_m$ is an annulus with one boundary component in $S^3 \times \{0\}$ and one in $S^3 \times \{1\}$. Then we say that the link $\partial_0\mathcal{W}$ is *order $n+1$ Whitney tower concordant* to $\partial_1\mathcal{W}$. This allows us to define the associated graded set W_n as $\mathbb{W}_n$ modulo order $n+1$ Whitney tower concordance. Knots have a well-defined connected sum operation, but the analogous band-sum operation for links is not well-defined, even up to concordance. This makes the following proposition somewhat surprising; it follows from Theorem 1.

Proposition 5. ([12]) *Band sum of links induces a well-defined operation which makes each W_n into a finitely generated abelian group.*

Our goal is to determine these groups W_n.

Free Lie and quasi-Lie algebras. 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, $\mathsf{L} = \oplus_n \mathsf{L}_n$, where the degree n part L_n is the additive abelian group of length n brackets, modulo Jacobi identities and the self-annihilation relations $[X,X]=0$. The free *quasi-Lie algebra* L' is gotten from L by replacing the self-annihilation relations with the weaker anti-symmetry relations $[X,Y] = -[Y,X]$.

The bracketing map $\mathsf{L}_1 \otimes \mathsf{L}_{n+1} \to \mathsf{L}_{n+2}$, has a nontrivial kernel, denoted D_n. The analogous bracketing map on the free quasi-Lie algebra is denoted D'_n. For later purposes, we now define a homomorphism $s\ell_{2n} : \mathsf{D}_{2n} \to \mathbb{Z}_2 \otimes \mathsf{L}_{n+1}$. Given an element X in D_{2n}, its image under the bracketing map is zero in L_{2n+2}. However, regarding the bracket as being in L'_{2n+2}, we get an element of the kernel of the projection $\mathsf{L}'_{2n+2} \to \mathsf{L}_{2n+2}$.

This kernel is isomorphic to $\mathbb{Z}_2 \otimes \mathsf{L}_{n+1}$ by [25], and so we get an element $s\ell_{2n}(X)$ of $\mathbb{Z}_2 \otimes \mathsf{L}_{n+1}$ as desired.

The total Milnor invariant. Let L be a link where all the longitudes lie in Γ_{n+1}, the $(n+1)$th term of the lower central series of the link group $\Gamma := \pi_1(S^3 \setminus L)$. As a consequence of Stallings' Theorem [31] it follows that $\frac{\Gamma_{n+1}}{\Gamma_{n+2}} \cong \frac{F_{n+1}}{F_{n+2}} \cong \mathsf{L}_{n+1}$, where $F = F(m)$ is the free group on meridians. Let $\mu_n^i(L) \in \mathsf{L}_{n+1}$ denote the image of the i-th longitude. The *total Milnor invariant* $\mu_n(L)$ of order n is defined by

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

It turns out that in fact $\mu_n(L) \in \mathsf{D}_n$ (by "cyclic symmetry"[27]). The invariant $\mu_n(L)$ is a convenient way of packaging all Milnor invariants of length $n+2$ in one piece.

Theorem 6. ([15]) *For all $n \in \mathbb{N}$, the total Milnor invariant is a well-defined homorphism $\mu_n : \mathsf{W}_n \to \mathsf{D}_n$ such that*

(i) For even n, μ_n is a monomorphism with image $\mathsf{D}'_n < \mathsf{D}_n$.
(ii) For odd n, μ_n is an epimorphism; denote its kernel by K^μ_n.

So μ_n is an algebraic obstruction for L bounding a Whitney tower of order $n+1$ which is a complete invariant in half the cases. In the other half, we'll need the following additional invariants.

Higher-order Sato-Levine invariants. Suppose $L \in \mathbb{W}_{2n-1}$ represents an element of K^μ_{2n-1}. Since $\mu_{2n-1}(L) = 0$, the longitudes lie in Γ_{2n}, so $\mu_{2n}(L) \in \mathsf{D}_{2n}$ is defined. Define the *order* $2n-1$ *Sato-Levine invariant* by $\mathrm{SL}_{2n-1}(L) = s\ell_{2n} \circ \mu_{2n}(L)$, where $s\ell_{2n}$ is defined above.

Theorem 7. ([15]) *For all n, the Sato-Levine invariant gives a well-defined epimorphism $\mathrm{SL}_{2n-1} : \mathsf{K}^\mu_{2n-1} \twoheadrightarrow \mathbb{Z}_2 \otimes \mathsf{L}_{n+1}$. Moreover, it is an isomorphism for even n.*

The case SL_1 is the original Sato-Levine [28] invariant of a 2-component classical link, and we describe in [15] (and below) how the SL_{2n-1} are obstructions to "untwisting" an order $2n$ twisted Whitney tower.

Higher-order Arf invariants. We saw above that the structure of the groups W_n is completely determined for $n \equiv 0, 2, 3$ mod 4 by Milnor and higher-order Sato-Levine invariants.

Theorem 8. ([15]) *Let $\mathsf{K}^{\mathrm{SL}}_{4k-3}$ be the kernel of SL_{4k-3}. Then there is an epimorphism $\alpha_k : \mathbb{Z}_2 \otimes \mathsf{L}_k \twoheadrightarrow \mathsf{K}^{\mathrm{SL}}_{4k-3}$.*

Conjecture 9. *α_k is an isomorphism.*

This conjecture is true when $k = 1$, and indeed the inverse map $\alpha_1^{-1} : \mathsf{W}_1 \to \mathbb{Z}_2 \otimes \mathsf{L}_1$ is given by the classical Arf invariant of each component of the link.

Regardless of whether or not Conjecture 9 is true, α_k induces an isomorphism $\overline{\alpha}_k$ on $(\mathbb{Z}_2 \otimes \mathsf{L}_k)/\operatorname{Ker}\alpha_k$.

Definition 2. *The* higher-order Arf invariants *are defined by*

$$\mathrm{Arf}_k := (\overline{\alpha}_k)^{-1} : \mathsf{K}^{\mathrm{SL}}_{4k-3} \to (\mathbb{Z}_2 \otimes \mathsf{L}_k)/\operatorname{Ker}\alpha_k$$

Any of the Arf_k which are non-trivial would the only possible remaining obstructions to a link bounding a Whitney tower of order $4k-2$, following the Milnor and Sato-Levine invariants:

Corollary 10. ([15]) *The associated graded groups W_n are classified by μ_n, SL_n if n is odd, and, for $n = 4k-3$, Arf_k.*

The first unknown Arf invariant is $\mathrm{Arf}_2 : \mathsf{W}_5 \to \mathbb{Z}_2 \otimes \mathsf{L}_2$, which in the case of 2-component links would be a $\mathbb{Z}_2$-valued invariant, evaluating non-trivially on the Bing double of any knot with non-trivial classical Arf invariant. Evidence supporting the existence of non-trivial Arf_k is provided by the fact that such links are known to not be slice [29]. All cases for $k > 1$ are currently unknown, but if Arf_2 is trivial then all higher-order Arf_k would also be trivial [13].

Twisted Whitney towers

The order n Sato-Levine invariants are defined as a certain projection of order $n+1$ Milnor invariants, suggesting that a slightly modified version of the Whitney tower filtration would put the Milnor invariants all in the right order, with no more need for the Sato-Levine invariants. In this section we discuss how this corresponds to the geometric notion of *twisted* Whitney towers.

Twisted Whitney disks. The normal disk-bundle of a Whitney disk $W \looparrowright M$ is isomorphic to $D^2 \times D^2$, and comes equipped with a canonical nowhere-vanishing *Whitney section* over the boundary given by pushing ∂W tangentially along one sheet and normally along the other.

The Whitney section determines the relative Euler number $\omega(W) \in \mathbb{Z}$ which represents the obstruction to extending the Whitney section across W. It depends only on a choice of orientation of the tangent bundle of the ambient 4-manifold restricted to the Whitney disk, i.e. a local orientation. 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.) If $\omega(W) = k$, we say that W is *k-twisted*, or just *twisted* if the value of $\omega(W)$ is not specified.

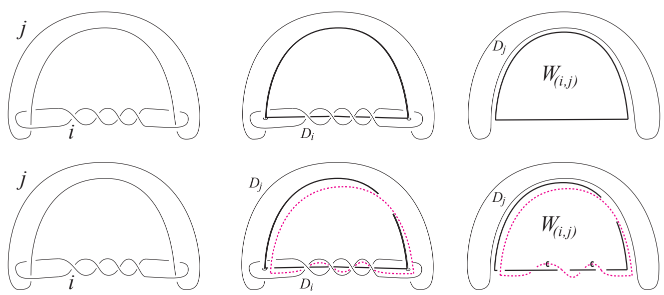

Fig. 4. Pushing into the 4–ball from left to right: An i- and j-labeled twisted Bing double of the unknot bounds disks D_i and D_j, which support a 2-twisted Whitney disk $W_{(i,j)}$. The Whitney section is indicated by the dotted red loop in the bottom center, and the intersections between its extension and the Whitney disk are shown in the bottom right.

In the definition of an order n Whitney tower given above all Whitney disks are required to be framed (0-twisted). It turns out that the natural generalization to twisted Whitney towers involves allowing non-trivially twisted Whitney disks only in at least "half the order" as follows:

Definition 3. *A* twisted Whitney tower of order $(2n-1)$ *is just a (framed) Whitney tower of order $(2n-1)$ as in Definition 1 above.*

A twisted Whitney tower of order $2n$ *is a Whitney tower having all intersections of order less than $2n$ paired by Whitney disks, with all Whitney disks of order less than n required to be framed, but Whitney disks of order at least n allowed to be k-twisted for any k.*

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 3.

Twisted Whitney tower concordance. Let $\mathbb{W}^{\infty}_n$ be the set of framed links in S^3 which are boundaries of order n twisted Whitney towers in B^4, with no requirement that the link framing is induced by the order 0 disks. Notice that $\mathbb{W}^{\infty}_{2n-1} = \mathbb{W}_{2n-1}$. While not immediately obvious, it is true that this defines a filtration $\cdots \subseteq \mathbb{W}^{\infty}_3 \subseteq \mathbb{W}^{\infty}_2 \subseteq \mathbb{W}^{\infty}_1 \subseteq \mathbb{W}^{\infty}_0 = \mathbb{L}$. As in the framed setting above, letting W^{∞}_n be the set $\mathbb{W}^{\infty}_n$ modulo order $(n+1)$ twisted Whitney tower concordance yields a finitely generated abelian group.

Theorem 11. ([13, 15]) *The total Milnor invariants give epimorphisms $\mu_n\colon \mathsf{W}^{\infty}_n \twoheadrightarrow \mathsf{D}_n$ which are isomorphisms for $n \equiv 0, 1, 3 \mod 4$. Moreover, the kernel $\mathsf{K}^{\infty}_{4k-2}$ of μ_{4k-2} is isomorphic to the kernel $\mathsf{K}^{\mathrm{SL}}_{4k-3}$ of the Sato-Levine map from the previous section.*

Conjecture 9 hence says that $\mathsf{K}^{\infty}_{4k-2} \cong \mathbb{Z}_2 \otimes \mathsf{L}_k$ and our Arf-invariants Arf_k represent the only remaining obstruction to a link bounding an order $4k-1$ twisted Whitney tower:

Corollary 12. *The groups W^{∞}_n are classified by μ_n and, for $n = 4k-2$, Arf_k.*

Gropes and k-slice links. Roughly speaking, a link is said to be "k-slice" if it is the boundary of a surface which "looks like a collection of slice disks modulo k-fold commutators in the fundamental group of the complement of the surface". Precisely, $L \subset S^3$ is *k-slice* if L bounds an embedded orientable surface $\Sigma \subset B^4$ such that $\pi_0(L) \to \pi_0(\Sigma)$ is a bijection and there is a push-off homomorphism $\pi_1(\Sigma) \to \pi_1(B^4 \setminus \Sigma)$ whose image lies in the kth term of the lower central series $(\pi_1(B^4 \setminus \Sigma))_k$. Igusa and Orr proved the following "k-slice conjecture" in [30]:

Theorem 13. ([30]) *A link L is k-slice if and only if $\mu_i(L) = 0$ for all $i \leq 2k-2$.*

A k-fold commutator in $\pi_1 X$ has a nice topological model in terms of a continuous map $G \to X$, where G is a *grope of class k*. Such 2-complexes G (with specified "boundary" circle) are recursively defined as follows. A grope of class 1 is a circle. A grope of class 2 is an orientable surface with one boundary component. A grope of class k is formed by attaching to every dual pair of basis curves on a class 2 grope a pair of gropes whose classes add to k. A curve $\gamma\colon S^1 \to X$ in a topological space X is a k-fold commutator if and only if it extends to a continuous map of a grope of class k. Thus one can ask whether being k-slice implies there is a basis of curves on Σ that bound *disjointly embedded* gropes of class k in $B^4 \setminus \Sigma$. Call such a link *geometrically k-slice.*

Proposition 14. ([13]) *A link L is geometrically k-slice if and only if $L \in \mathbb{W}^{\infty}_{2k-1}$.*

This is proven using a construction from [17] which allows one to freely pass between class n gropes and order $n-1$ Whitney towers. So the higher-order Arf invariants Arf_k detect the difference between k-sliceness and geometric k-sliceness. It turns out that every Arf_k value can be realized by (internal) band summing iterated Bing doubles of the figure-eight knot. Every Bing double is a boundary link, and one can choose the bands so that the sum remains a boundary link. This implies:

Theorem 15. ([13]) *A link L has vanishing Milnor invariants of all orders $\leq 2k-2$ if and only if it is geometrically k-slice after connected sums with internal band sums of iterated Bing doubles of the figure-eight knot.*

Here (and in Theorem 17 below), the figure-eight knot can be replaced by any knot with non-trivial (classical) Arf invariant.

The added boundary links in the above theorem bound disjoint surfaces in S^3 which clearly allow immersed disks in B^4 bounded by curves representing a basis of first homology. In [13] we will show that this implies:

Theorem 16. ([13]) *A link has vanishing Milnor invariants of all orders $\leq 2k-2$ if and only if its components bound disjointly embedded surfaces $\Sigma_i \subset B^4$, with each surface a connected sum of two surfaces Σ'_i and Σ''_i such that*

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

This is an enormous geometric strengthening of Igusa and Orr's result, which under the same assumption on the vanishing of Milnor invariants, shows the existence of a surface Σ with a basis of curves bounding *maps* of class k gropes, with no control on their intersections and self-intersections. Our proof uses the full power of the obstruction theory for twisted Whitney towers, whereas they do a sophisticated computation of the third homology of the groups F/F_{2k}.

String links and the Artin representation. Let L be a string link with m strands embedded in $D^2 \times [0,1]$. By Stallings' Theorem [31], the inclusions $(D^2 \setminus \{m \text{ points}\}) \times \{i\} \hookrightarrow (D^2 \times [0,1]) \setminus L$ for $i = 0, 1$ induce isomorphisms on all lower central quotients of the fundamental groups. In fact, the induced automorphism of the lower central quotients F/F_n of the free group $F = \pi_1(D^2 \setminus \{m \text{ points}\})$ is explicitly characterized by conjugating the meridional generators of F by longitudes. Let $\mathrm{Aut}_0(F/F_n)$ consist of those automorphisms of F/F_n which are defined by conjugating each generator and which fix the product of generators. This leads to the *Artin representation* $\mathbb{SL} \to \mathrm{Aut}_0(F/F_{n+2})$ where $\mathbb{SL}$ is the set of concordance classes of pure framed string links.

The set of string links has an advantage over links in that it has a well-defined monoid structure given by stacking. Indeed, modulo concordance, it becomes a (noncommutative) group. Whitney tower filtrations can also be defined in this context, giving rise to filtrations $\mathbb{SW}_n$ and $\mathbb{SW}^{\infty}_n$ of this group $\mathbb{SL}$.

Theorem 17. ([16]) *The sets $\mathbb{SW}_n$ and $\mathbb{SW}^{\infty}_n$ are normal subgroups of $\mathbb{SL}$ which are central modulo the next order. We obtain nilpotent groups $\mathbb{SL}/\mathbb{SW}_n$ and $\mathbb{SL}/\mathbb{SW}^{\infty}_n$, and the associated graded groups are isomorphic to our previously defined groups:*

$$\mathbb{SW}_n/\mathbb{SW}_{n+1} \cong \mathsf{W}_n \quad \textit{and} \quad \mathbb{SW}^{\infty}_n/\mathbb{SW}^{\infty}_{n+1} \cong \mathsf{W}^{\infty}_n$$

Finally, the Artin representation induces a well-defined epimorphism $\mathrm{Artin}_n\colon \mathbb{SL}/\mathbb{SW}^{\infty}_n \twoheadrightarrow \mathrm{Aut}_0(F/F_{n+2})$ whose kernel is generated by internal band sums of iterated Bing doubles of the figure-eight knot.

The Artin representation is thus an invariant on the whole group $\mathbb{SL}/\mathbb{SW}^{\infty}_n$, not just on the associated graded groups as in the case of links. It packages the total Milnor invariants μ_k, $k = 0, \ldots, n$ on string links together into a group homomorphism. (See [16] for Bing-doubling string links.)

Higher-order intersection invariants

Proofs of the above results depend on two essential ideas: The higher-order intersection theory of Whitney towers comes with an obstruction theory whose associated invariants take values in abelian groups of (un-rooted) trivalent trees. And by mapping to rooted trees, which correspond to iterated commutators, the obstruction theory for Whitney towers in B^4 can be identified with algebraic invariants of the bounding link in S^3. A critical connection between these ideas is provided by the resolution of the Levine Conjecture (see below), which says that this map is an isomorphism.

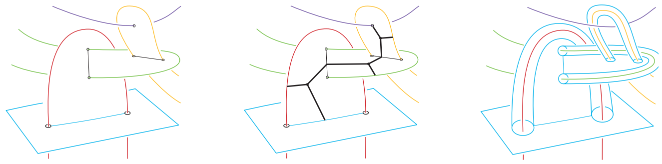

Fig. 5. From left to right: An unpaired intersection in a Whitney tower, (part of) its associated tree, and the result of surgering to a grope.

In fact, it can be arranged that all singularities in a Whitney tower are contained in 4-ball neighborhoods of the associated trivalent trees, which sit as embedded 'spines'; and all relations among trees in the target group are realized by controlled manipulations of the Whitney disks. Mapping to rooted trees corresponds geometrically to surgering Whitney towers to gropes, and these determine iterated commutators of meridians of the Whitney tower boundaries as in Figure 5.

Trees and intersections. All trees are unitrivalent, with cyclic orderings of the edges at all trivalent vertices, and univalent vertices labeled from an index set $\{1, 2, 3, \ldots, m\}$. A *rooted* tree has one unlabeled univalent vertex designated as the *root*. Such rooted trees correspond to formal non-associative bracketings of elements from the index set. The *rooted product* (I, J) of rooted trees I and 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, and we identify rooted trees with formal brackets. The *inner product* $\langle I, J\rangle$ of rooted trees I and J is the unrooted tree gotten by identifying the roots of I and J to a single non-vertex point. Note that all the univalent vertices of $\langle I, J\rangle$ are labeled.

The *order* of a tree, rooted or unrooted, is defined to be the number of trivalent vertices, and the following associations of trees to Whitney disks and intersection points respects the notion of order given in Definition 1.

To each order zero surface A_i is associated the order zero rooted tree consisting of an edge with one vertex labeled by i, and to each transverse intersection $p \in A_i \cap A_j$ is associated the order zero tree $t_p := \langle i, j\rangle$ consisting of an edge with vertices labeled by i and j. 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 an embedded subset of 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.

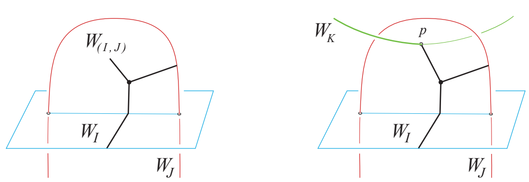

Fig. 6.

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 6); with the understanding that if, say, I is just a singleton i, then W_I denotes the order zero 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 un-rooted tree $t_p := \langle (I, J), K\rangle$ (see right-hand side of Figure 6).

Intersections trees for Whitney towers. The group $\mathcal{T}_n$ (for each $n = 0, 1, 2 \ldots$) is the free abelian group on (unitrivalent labeled vertex-oriented) order n trees, modulo the usual AS (antisymmetry) and IHX (Jacobi) relations:

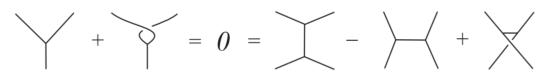

In even orders we define $\widetilde{\mathcal{T}}_{2n} := \mathcal{T}_{2n}$, and in odd orders $\widetilde{\mathcal{T}}_{2n-1}$ is defined to be the quotient of $\mathcal{T}_{2n-1}$ by the *framing relations.* These framing relations are defined as the image of homomorphisms $\Delta_{2n-1} : \mathbb{Z}_2 \otimes \mathcal{T}_{n-1} \to \mathcal{T}_{2n-1}$ which are defined for generators $t \in \mathcal{T}_{n-1}$ by $\Delta(t) := \sum_{v\in t}\langle i(v), (T_v(t), T_v(t))\rangle$, where $T_v(t)$ denotes the rooted tree gotten by replacing v with a root, and the sum is over all univalent vertices of t, with $i(v)$ the original label of the univalent vertex v.

The obstruction theory works as follows:

Definition 4. *The* order n intersection tree $\tau_n(\mathcal{W})$ *of an order n Whitney tower $\mathcal{W}$ is defined to be*

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

where the sum is over all order n intersections p, with $\epsilon_p = \pm 1$ the usual sign of a transverse intersection point (via certain orientation conventions, see e.g. [12]).

All relations in $\widetilde{\mathcal{T}}_n$ can be realized by controlled manipulations of Whitney towers, and further maneuvers allow algebraically canceling pairs of tree generators to be converted into intersection-point pairs admitting Whitney disks. As a result, we get the following partial recovery of the "algebraic cancelation implies geometric cancellation" principle available in higher dimensions:

Theorem 18. ([12]) *If a collection A of properly immersed surfaces in a simply-connected 4–manifold supports an order n Whitney tower $\mathcal{W}$ with $\tau_n(\mathcal{W}) = 0 \in \widetilde{\mathcal{T}}_n$, then A is homotopic (rel ∂) to A' which supports an order $n+1$ Whitney tower.*

Intersections trees for twisted Whitney towers. For any rooted tree J we define the corresponding ∞*-tree* ("twisted-tree"), denoted by J^∞, by labeling the root univalent vertex with the symbol "∞" (which will represent a "twist" in a Whitney disk normal bundle): $J^\infty := \infty\!-\!\!-J$.

Definition 5. *The group $\mathcal{T}^\infty_{2n-1}$ is the quotient of $\widetilde{\mathcal{T}}_{2n-1}$ by the* boundary-twist relations*:*

$$\langle (i, J), J\rangle \;=\; i \prec^{J}_{J} \;=\; 0$$

Here J ranges over all order $n-1$ rooted trees (and the first equality is just a reminder of notation).

The group $\mathcal{T}^\infty_{2n}$ is gotten from $\widetilde{\mathcal{T}}_{2n} = \mathcal{T}_{2n}$ by including order n ∞-trees as new generators and introducing the following new relations (in addition to the IHX and antisymmetry relations on non-∞ trees):

$$J^\infty = (-J)^\infty \qquad I^\infty = H^\infty + X^\infty - \langle H, X\rangle \qquad 2\cdot J^\infty = \langle J, J\rangle$$

The left-hand *symmetry* relation corresponds to the fact that the framing obstruction on a Whitney disk is independent of its orientation; the middle *twisted IHX* relations can be realized by a Whitney move near a twisted Whitney disk, and the right-hand *interior twist* relations can be realized by cusp-homotopies in Whitney disk interiors. As described in [15], the twisted groups $\mathcal{T}^\infty_{2n}$ can naturally be identified with a universal quadratic refinement of the $\mathcal{T}_{2n}$-valued intersection pairing $\langle\cdot,\cdot\rangle$ on framed Whitney disks.

Recalling from Definition 3 that twisted Whitney disks only occur in even order twisted Whitney towers, intersection trees for twisted Whitney towers are defined as follows:

Definition 6. *The* order n intersection tree $\tau_n^\infty(\mathcal{W})$ *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 $\omega(W_J) \in \mathbb{Z}$ (computed from a consistent choice of local orientations).

By "splitting" the twisted Whitney disks [12] it can be arranged that $|\omega(W_J)| \leq 1$, leading to signs like ϵ_p (or zero coefficients). The obstruction theory also holds for twisted Whitney towers:

Theorem 19. ([12]) *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 homotopic (rel ∂) to A' which supports an order $n+1$ twisted Whitney tower.*

Remark on the framing relations. The framing relations in the groups $\widetilde{\mathcal{T}}_{2n-1}$ correspond to the twisted IHX relations among ∞-trees in $\mathcal{T}_{2n}^\infty$ via a geometric boundary-twist operation which converts an order n ∞-tree $(i, J)^\infty$ to an order $2n-1$ (non-∞) tree $\langle (i, J), J \rangle$.

Realization maps. In [12] we describe how to construct surjective realization maps $\widetilde{R}_n : \widetilde{\mathcal{T}}_n \twoheadrightarrow \mathsf{W}_n$ and $R_n^\infty : \mathcal{T}_n^\infty \twoheadrightarrow \mathsf{W}_n^\infty$ by applying the operation of iterated Bing doubling. This construction is essentially the same as Cochran's realization method for Milnor invariants [32, 33] and Habiro's clasper-surgery [34], extended to twisted Bing doubling (Figures 7 and 4). To prove the realization maps are well-defined, we need to use Theorems 18 and 19 respectively.

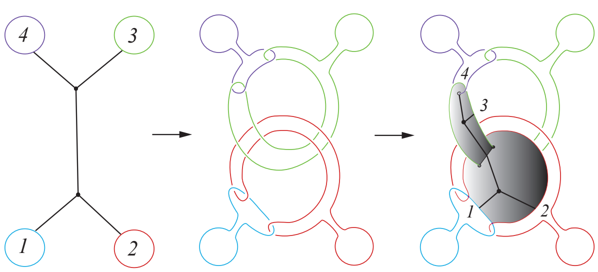

Fig. 7. Realizing an order 2 tree in a Whitney tower by Bing doubling.

The above Conjecture 9 on the non-triviality of the higher-order Arf invariants can be succinctly rephrased as the assertion that the realization maps $\widetilde{R}_n$ and R_n^∞ are *isomorphisms* for all n. Progress towards confirming this assertion – namely complete answers in 3/4 of the cases and partial answers in the remaining cases, as described by the above-stated results – has been accomplished by identifying intersection trees with Milnor invariants, as we describe next.

Intersection trees and Milnor's link invariants. The connection between intersection trees and Milnor invariants is via a surjective map $\eta_n : \mathcal{T}_n^\infty \to \mathsf{D}_n$ which converts trees to rooted trees (interpreted as Lie brackets) by summing over all ways of choosing a root:

For v a univalent vertex of an order n (un-rooted non-∞) tree 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 of indices 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 lies in $\mathsf{L}_1 \otimes \mathsf{L}_{n+1}$ because the coefficient of $\eta_n(\langle J, J \rangle)$ is even.

The proof of the following theorem (which implies Theorem 11 above) shows that the map η corresponds to a construction which converts Whitney towers into embedded gropes [17], via the grope duality of [35]:

Theorem 20. ([13]) *If L bounds a twisted Whitney tower $\mathcal{W}$ of order n, then the total Milnor invariants $\mu_k(L)$ vanish for $k < n$, and $\mu_n(L) = \eta_n \circ \tau_n^\infty(\mathcal{W}) \in \mathsf{D}_n$.*

Thus one needs to understand the kernel of η_n before the obstruction theory can proceed. This is accomplished by resolving [14] a closely related conjecture of J. Levine [24], as discussed next.

The Levine Conjecture and its implications. The bracket map kernel D_n turns out to be relevant to a variety of topological settings (see e.g. the introduction to [14]), and was known to be isomorphic to $\mathcal{T}_n$ after tensoring with $\mathbb{Q}$, when Levine's study of the cobordism groups of 3-dimensional homology cylinders [23, 24] led him to conjecture that $\mathcal{T}_n$ is in fact isomorphic to the quasi-Lie bracket map kernel D'_n, via the analogous map η'_n which sums over all choices of roots (as in the left formula for η above).

Levine made progress in [24, 25], and in [14] we affirm his conjecture:

Theorem 21. ([14]) *$\eta'_n : \mathcal{T}_n \to \mathsf{D}'_n$ is an isomorphism for all n.*

The proof of Theorem 21 uses techniques from discrete Morse theory on chain complexes, including an extension of the theory to complexes containing torsion. A key idea involves defining combinatorial vector fields that are inspired by the Hall basis algorithm for free Lie algebras and its generalization by Levine to quasi-Lie algebras.

As described in [15], Theorem 21 has several direct applications to Whitney towers, including the completion of the calculation of W_n^∞ in three out of four cases:

Theorem 22. ([15]) *$\eta_n : \mathcal{T}_n^\infty \to \mathsf{D}_n$ are isomorphisms for $n \equiv 0, 1, 3 \mod 4$. As a consequence, both the total Milnor invariants $\mu_n : \mathsf{W}_n^\infty \to \mathsf{D}_n$ and the realization maps $R_n^\infty : \mathcal{T}_n^\infty \to \mathsf{W}_n^\infty$ are isomorphisms for these orders.*

The consequences listed in the second statement follow from the fact that η_n is the composition

$$\eta_n : \mathcal{T}_n^\infty \xrightarrow{R_n^\infty} \mathsf{W}_n^\infty \xrightarrow{\mu_n} \mathsf{D}_n$$

Theorem 21 is also instrumental in determining the only possible remaining obstructions to computing W_{4k-2}^∞:

Proposition 23. ([15]) *The map sending a rooted tree J to $(J, J)^\infty \in \mathcal{T}_{4k-2}^\infty$ induces an isomorphism*

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

These symmetric ∞-trees $(J, J)^\infty$ correspond to twisted Whitney disks, and determine the higher-order Arf invariants Arf_k. All of our above conjectures are equivalent to the statement that W_{4k-2}^∞ is isomorphic to $\mathsf{D}_{4k-2} \oplus (\mathbb{Z}_2 \otimes \mathsf{L}_k)$ via these maps.

Theorem 22 and Proposition 23 imply Theorem 11 and Corollary 12 above, and [15] describes analogous implications of the above-described results in the framed setting (Theorems 6, 7, 8, and Corollary 10).

Framed versus twisted Whitney towers

This section describes how the higher-order Sato-Levine and Arf invariants can be interpreted as obstructions to framing a twisted Whitney tower. The starting point is the following surprisingly simple relation between twisted and framed Whitney towers of various orders:

Proposition 24. ([12, 15]) *For any $n \in \mathbb{N}$, there is a commutative diagram of exact sequences*

$$\begin{tikzcd}
0 \arrow[r] & \mathsf{W}_{2n} \arrow[r] & \mathsf{W}^{\infty}_{2n} \arrow[r] & \mathsf{W}_{2n-1} \arrow[r] & \mathsf{W}^{\infty}_{2n-1} \arrow[r] & 0 \\
0 \arrow[r] & \widetilde{\mathcal{T}}_{2n} \arrow[r] \arrow[u, two heads, "\widetilde{R}_{2n}"] & \mathcal{T}^{\infty}_{2n} \arrow[r] \arrow[u, two heads, "R^{\infty}_{2n}"] & \widetilde{\mathcal{T}}_{2n-1} \arrow[r] \arrow[u, two heads, "\widetilde{R}_{2n-1}"] & \mathcal{T}^{\infty}_{2n-1} \arrow[r] \arrow[u, two heads, "R^{\infty}_{2n-1}"] & 0
\end{tikzcd}$$

Moreover, there are isomorphisms

$$\mathrm{Cok}(\mathcal{T}_{2n} \to \mathcal{T}^{\infty}_{2n}) \cong \mathbb{Z}_2 \otimes \mathsf{L}'_{n+1} \cong \mathrm{Ker}(\widetilde{\mathcal{T}}_{2n-1} \to \mathcal{T}^{\infty}_{2n-1})$$

In the first row, all maps are induced by the identity on the set of links. To see the exactness, observe that there is a natural inclusion $\mathbb{W}_n \subseteq \mathbb{W}^{\infty}_n$, and by definition $\mathbb{W}^{\infty}_{2n-1} = \mathbb{W}_{2n-1}$. One then needs to show that indeed $\mathbb{W}^{\infty}_{2n} \subseteq \mathbb{W}^{\infty}_{2n-1}$, which is accomplished in [12], and then the exact sequence in Proposition 24 follows since $\mathsf{W}_n := \mathbb{W}_n/\mathbb{W}_{n+1}$ and $\mathsf{W}^{\infty}_n := \mathbb{W}^{\infty}_n/\mathbb{W}^{\infty}_{n+1}$.

If our above conjectures hold, then for every n the various (vertical) realization maps in the above diagram are isomorphisms, which would lead to a computation of the cokernel and kernel of the map $\mathsf{W}_n \to \mathsf{W}^{\infty}_n$. As a consequence, we would obtain new concordance invariants with values in $\mathbb{Z}_2 \otimes \mathsf{L}'_{n+1}$ and defined on $\mathbb{W}^{\infty}_{2n}$, as the obstructions for a link to bound a framed Whitney tower of order $2n$. In fact [15], the above-defined higher-order Sato-Levine invariants detect the quotient $\mathbb{Z}_2 \otimes \mathsf{L}_{n+1}$ of $\mathbb{Z}_2 \otimes \mathsf{L}'_{n+1}$. Levine [24] showed that the squaring map $X \mapsto [X, X]$ induces an isomorphism

$$\mathbb{Z}_2 \otimes \mathsf{L}_k \cong \mathrm{Ker}(\mathbb{Z}_2 \otimes \mathsf{L}'_{2k} \twoheadrightarrow \mathbb{Z}_2 \otimes \mathsf{L}_{2k}),$$

which leads to our proposed higher-order Arf invariants Arf_k.

It is interesting to note that the case $n = 0$ leads to the prediction $\mathrm{Cok}(\mathsf{W}_0 \to \mathsf{W}^{\infty}_0) \cong \mathbb{Z}_2 \otimes \mathsf{L}_1 \cong (\mathbb{Z}_2)^m$ This is indeed the group of framed m-component links modulo those with even framings! In fact, the consistency of this computation was the motivating factor to consider filtrations of the set of *framed* links $\mathbb{L}$, rather than just oriented links.

Filtrations of homology cylinders

Garoufalidis and Levine [36] studied the group $\mathcal{H}_g$ of *homology cylinders* over the compact orientable surface of genus g with one boundary component, modulo *homology cobordism.* It carries the Johnson (relative weight) filtration $\mathbb{J}_n$ and the Goussarov-Habiro (clasper) filtration $\mathbb{Y}_n$. We improve results on the comparison of the associated graded groups J_n and Y_n.

Theorem 25. ([16]) *For all $k \geq 1$, there are exact sequences*

(i) $0 \to \mathsf{Y}_{2k} \to \mathsf{J}_{2k} \to \mathbb{Z}_2 \otimes \mathsf{L}_{k+1} \to 0$
(ii) $0 \to \mathbb{Z}_2 \otimes \mathsf{L}_{2k+1} \to \mathsf{Y}_{4k-1} \to \mathsf{J}_{4k-1} \to 0$
(iii) $0 \to \mathsf{K}^{\mathsf{Y}}_{4k-3} \to \mathsf{Y}_{4k-3} \to \mathsf{J}_{4k-3} \to 0$
(iv) $\mathbb{Z}_2 \otimes \mathsf{L}_k \xrightarrow{a_k} \mathsf{K}^{\mathsf{Y}}_{4k-3} \to \mathbb{Z}_2 \otimes \mathsf{L}_{2k} \to 0.$

ACKNOWLEDGMENTS. This paper was written while the first two authors were visiting the third author at the Max-Planck-Institut für Mathematik in Bonn. They all thank MPIM for its stimulating research environment and generous support. The last author was also supported by NSF grants DMS-0806052 and DMS-0757312.

1. H Whitney, (1944) The self intersections of a smooth n–manifold in 2n–space, Annals of Math. 45 220–246.
2. S Smale, (1962) On the structure of manifolds Amer. J. Math. , 84 387–399
3. T Wall, (1970) Surgery on Compact Manifolds, London Math.Soc.Monographs 1, Academic Press or Second Edition, edited by A. Ranicki, Math. Surveys and Monographs 69, A.M.S.
4. A Casson, (1974) Three Lectures on new infinite constructions in 4–dimensional manifolds, Notes prepared by L Guillou, Prepublications Orsay 81T06.
5. M Freedman, (1983) The disk theorem for four-dimensional manifolds, Proc. ICM Warsaw, 647–663
6. M Freedman, (1982) The topology of four-dimensional manifolds, J. Diff. Geom. 17 357–453.
7. S Donaldson, (1983) Self-dual connections and the topology of smooth 4–manifolds, Bull. Amer. Math. Soc. 8, 81–83.
8. M Freedman, F Quinn, (1990) The topology of 4–manifolds, Princeton Math. Series 39 Princeton, NJ.
9. M Freedman, P Teichner, (1995) 4–manifold topology I: Subexponential groups, Invent. Math. 122 509–529.
10. V Krushkal, F Quinn, (2000) Subexponential groups in 4–manifold topology, Geometry and Topology Vol. 4 407–430.
11. J Conant, R Schneiderman, P Teichner, (2007) Jacobi identities in low-dimensional topology, Compositio Mathematica 143 Part 3 780–810.
12. J Conant, R Schneiderman, P Teichner, (2010) Geometric Filtrations of Classical Link Concordance, preprint http://arxiv.org/abs/1101.3477
13. J Conant, R Schneiderman, P Teichner, (2010) Milnor invariants and twisted Whitney towers, preprint http://arxiv.org/abs/1102.0758
14. J Conant, R Schneiderman, P Teichner, (2010) Tree homology and a conjecture of Levine, preprint http://arxiv.org/abs/1012.2780
15. J Conant, R Schneiderman, P Teichner, (2010) Universal quadratic refinements and untwisting Whitney towers, preprint http://arxiv.org/abs/1101.3480
16. J Conant, R Schneiderman, P Teichner, (2010) Whitney tower concordance, string links and homology cylinders, preprint.
17. R Schneiderman, (2006) Whitney towers and Gropes in 4–manifolds, Trans. Amer. Math. Soc. 358 4251–4278.
18. R Schneiderman, P Teichner, (2001) Higher order intersection numbers of 2–spheres in 4–manifolds, Alg. and Geom. Topology 1 1–29.
19. R Schneiderman, P Teichner, (2004) Whitney towers and the Kontsevich integral, Proceedings of the Casson Fest, Geometry and Topology Monograph Series.
20. T Cochran, K Orr, P Teichner, (2003) Knot concordance, Whitney towers and L^2-signatures, Annals of Math., Volume 157 433–519.
21. J Milnor, (1954) Link groups, Annals of Math. 59 177–195.
22. J Milnor, (1957) Isotopy of links, Algebraic geometry and topology, Princeton Univ. Press.
23. J Levine, (2001) Homology cylinders: an enlargement of the mapping class group, Alg. and Geom. Topology 1 243–270.
24. J Levine, (2002) Addendum and correction to: Homology cylinders: an enlargement of the mapping class group, Alg. and Geom. Topology 2 1197–1204.
25. J Levine, (2006) Labeled binary planar trees and quasi-Lie algebras, Alg. and Geom. Topology 6 935–948.
26. R Schneiderman, (2010) Stable concordance of knots in 3–manifolds, Alg. and Geom. Topology 10 373–432.
27. M Freedman, P Teichner, (1995) 4–manifold topology II: Dwyer's filtration and surgery kernels, Invent. Math. 122 531–557.
28. N Sato, (1984) Cobordisms of semi-boundary links, Topology Appl. 18 225–234.
29. J C Cha, (2010) Link concordance, homology cobordism, and Hirzebruch-type defects from iterated p -covers, Journal of the European Mathematical Society 12, 555–610.
30. K Igusa, K Orr, (2001) Links, pictures and the homology of nilpotent groups, Topology 40 1125–1166.
31. J Stallings, (1965) Homotopy and central series of groups, J. Algebra, 170–181.
32. T Cochran, (1990) Derivatives of links, Milnor's concordance invariants and Massey products, Mem. Amer. Math. Soc. Vol. 84 No. 427.
33. T Cochran, (1991) k-cobordism for links in S^3, Trans. AMS 327, No. 2, 641–654.
34. K Habiro, (2000) Claspers and finite type invariants of links, Geom. Topol. Vol. 4 1–83.
35. V Krushkal, P Teichner, (1997) Alexander duality, Gropes and link homotopy, Geom. Topol. Vol. 1 51–69.
36. S Garoufalidis, J Levine, (2001) Homology surgery and invariants of 3–manifolds, Geom. Topol. Vol. 5 75–108.