Journal of Knot Theory and Its Ramifications
Vol. 26, No. 2 (2017) 1740012 (28 pages)

DOI: 10.1142/S0218216517400120

World Scientific
www.worldscientific.com

Cochran's β^i-invariants via twisted Whitney towers

Jim Conant*, Rob Schneiderman† and Peter Teichner‡

Department of Mathematics,
University of Tennessee, Knoxville, USA

Department of Mathematics and Computer Science,
Lehman College, City University of New York, USA

Max-Planck-Institut für Mathematik, Bonn, Germany

University of California, Berkeley, USA
**jconant@math.utk.edu*
†robert.schneiderman@lehman.cuny.edu
‡teichner@mac.com

Received 15 February 2016
Accepted 30 June 2016
Published 29 November 2016

ABSTRACT

We show that Cochran's invariants $\beta^i(L)$ of a 2-component link L in the 3-sphere can be computed as intersection invariants of certain 2-complexes in the 4-ball with boundary L. These 2-complexes are special types of twisted Whitney towers, which we call *Cochran towers*, and which exhibit a new phenomenon: A Cochran tower of order $2k$ allows the computation of the β^i invariants for all $i \leq k$, i.e. simultaneous extraction of invariants from a Whitney tower at multiple orders. This is in contrast with the order n Milnor invariants (requiring order n Whitney towers) and consistent with Cochran's result that the $\beta^i(L)$ are integer lifts of certain Milnor invariants.

Keywords: Cochran invariants; clasper calculus; link concordance; Milnor invariants; Whitney towers.

Mathematics Subject Classification 2010: 57M25, 57M27

1. Introduction and Statement of Results

In 1954, John Milnor defined his μ-invariants of a link $L = (L_1, \ldots, L_m)$ in 3-space [18] by looking inductively at the terms in the lower central series of the link group $\pi_1(\mathbb{R}^3 \smallsetminus L)$, and comparing with the link group of the unlink. For example, the *order* 0 Milnor invariants are just the linking numbers $\mu_{ij}(L)$ between components L_i and L_j of the link L. Moreover, Milnor showed that μ_{123} detects the Borromean rings, a *Bing double* of the Hopf link, and that his higher-order invariants detect iterated Bing doublings of the Hopf link.

The Milnor invariants $\mu_{\mathcal{I}}(L)$ of *order* $n \geq 0$ are labeled by a multi-index $\mathcal{I} = \{i_1 i_2 \ldots i_{n+2}\}$ with $i_k \in \{1, \ldots, m\}$. They are integers, well-defined only modulo

*Corresponding author.

the gcd of the $\mu_{\mathcal{I}'}(L)$ for all proper subindices $\mathcal{I}' \subset \mathcal{I}$. For example, if L is a two component link with linking number $\mu_{12}(L) = 1$, then all the higher-order Milnor invariants of L are completely ill-defined. Nevertheless, Milnor invariants turn out to be a central tool in 3- and 4-dimensional topology, particularly because of their concordance invariance.

For example, if the topological surgery sequence is exact in dimension 4 for free fundamental groups, then the *Whitehead double* of any link L with trivial linking numbers is topologically slice. However, this last statement is currently only known to hold for links L with $\mu_{\mathcal{I}}(L) = 0$ for any multi-index $\mathcal{I} = \{i_1 i_2 \dots i_{n+2}\}$ in which at most one index appears more than once (and at most twice) in $\mathcal{I}$ [12].

In 1985, Tim Cochran discovered a beautiful method of lifting certain Milnor invariants to well-defined integers [2]: Given a 2-component link $L = (L_1, L_2)$ with $\mu_{12}(L) = 0$, he first defined its *derived link* $D(L)$ by forming a knot as the intersection of Seifert surfaces for the components (each in the complement of the other component), and then taking this knot in place of L_2 to yield the new 2-component link $D(L)$.

Cochran then defined $\beta^1(L) \in \mathbb{Z}$ as the *Sato–Levine invariant* of L [19], which is the twisting of the intersection knot, inherited from either Seifert surface. What are today known as *higher-order Cochran invariants* were defined recursively via the formula $\beta^i(L) = \beta^{i-1}(D(L)) \in \mathbb{Z}$. Amazingly, these integers are well-defined for $i \geq 1$, and Cochran showed [3, Theorem 6.10] that they are lifts of the following Milnor invariants:

$$\beta^i(L) \equiv \mu_{1^{2^i}2^2}(L) := \mu_{1\dots 2_i \dots 122}(L) \mod \gcd\{\mu_{1^{2^k}2^2}(L), k < i\}$$

Only $\beta^1(L) = \mu_{1122}(L)$ was known to be a well-defined integer for linking number zero links.

In this paper, we re-interpret Cochran's invariants $\beta^i(L)$ in terms of intersection invariants of certain 2-complexes in the 4-ball B^4, with boundary $L \subset S^3$. These 2-complexes are special kinds of *twisted Whitney towers* [5, 6] and we propose to call them *Cochran towers*.

The notions necessary to understand our theorem below will be given in the remainder of the introduction but first we would like to state the main result:

Theorem 1.1. *For any $k \geq 1$, a link $L = (L_1, L_2)$ in S^3 bounds a Cochran tower of order $2k$ in B^4 if and only if L has trivial linking number and the Arf invariant of L_1 vanishes. Moreover, given a Cochran tower C of order $2k$ with boundary L then for any $i \leq k$, the Cochran invariants can be computed as follows*:

$$\beta^i(L) = \sum \omega(W_J),$$

where the sum is over all Whitney disks W_J in C with $J^{\infty} \cong t_i^{\infty}$ and the twisting $\omega(W_J) \in \mathbb{Z}$ is the relative Euler number of W_J.

Here J^{∞} is a trivalent tree associated to the Whitney disk W_J in C and the relevant tree t_i^{∞} is shown on the right-hand side of Fig. 1.

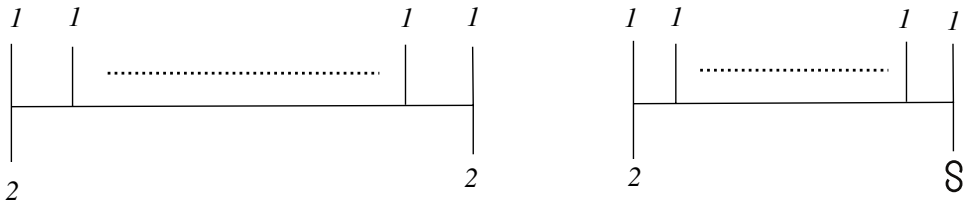

Fig. 1. The tree t_i on the left has two vertices labeled 2 and i vertices labeled 1. The tree t_i^∞ on the right has i vertices labeled 1, one vertex labeled 2, and one vertex labeled by the *twist* symbol.

Remark 1.2. If the ($\mathbb{Z}/2$-valued) Arf invariant $\mathrm{Arf}(L_1) \neq 0$, one can change L by tying a small trefoil knot into L_1. This does not alter β^1, nor the derived link $D(L)$ (and hence all $\beta^i(L)$ are unchanged) but allows one to build a Cochran tower on this new link to compute all $\beta^i(L)$.

A geometric interpretation similar to Theorem 1.1 was given in our earlier work [6] for the *first non-vanishing* Milnor invariants $\mu_{\mathcal{I}}(L) \in \mathbb{Z}$ of a given link L. This means that $\mu_{\mathcal{I}'}(L) = 0$ for all proper subsets $\mathcal{I}' \subset \mathcal{I}$. These interpretations use our theory of *twisted Whitney towers* in the 4-ball with boundary L, which is surveyed in [4] and detailed in [5, 6]. We will sketch next those aspects which are relevant to the current discussion, and point the reader to the appropriate references for more information. The smooth category is assumed throughout, except when otherwise specified.

1.1. *Twisted Whitney towers and their intersection forests*

Roughly speaking, a twisted Whitney tower (with boundary a link $L \subset S^3$) is a finite 2-complex $\mathcal{W} \subset B^4$ with $\mathcal{W} \cap S^3 = L$, which is the union of a finite number of transverse disks with carefully chosen boundary conditions. More precisely, $\mathcal{W}$ is formed by taking generic disks W_j in the 4-ball with $\partial W_j = L_j$ and then adding Whitney disks for pairs of intersection or self-intersection points among the W_j (if possible). One continues to inductively add *higher-order* Whitney disks for pairs of (self)-intersection points among previously included disks, to arrive at a Whitney tower after finitely many steps.

The main use of a Whitney tower $\mathcal{W}$ comes from its *intersection forest* $t(\mathcal{W})$ [6, Sec. 2.5] which is a disjoint union of trivalent labeled trees, as we review next. As mentioned after Definition 1.3, $t(\mathcal{W})$ represents obstructions to successfully carrying out Whitney moves that would lead to slice disks for the link L.

To define $t(\mathcal{W})$, start by associating a rooted unitrivalent tree J to each Whitney disk $W \subset \mathcal{W}$ as follows: Take a univalent *root* vertex of J sitting in the interior of W along with an edge to the adjacent trivalent vertex, and take the other edges of J to be sheet-changing paths bifurcating down through the lower-order Whitney disks until arriving at the boundary components L_j. This yields j-labeled univalent vertices in J (the root is the only unlabeled univalent vertex).

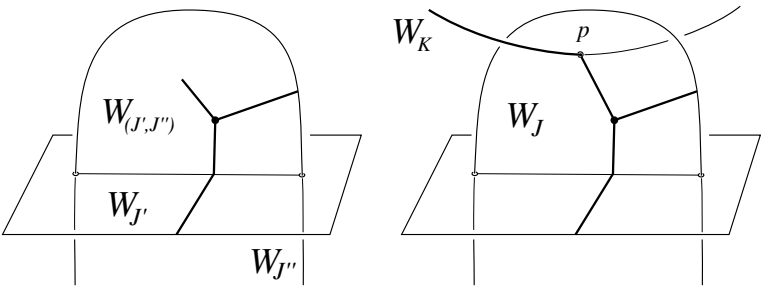

Fig. 2. Left: A Whitney disk $W_{(J',J'')}$ pairing intersections between $W_{J'}$ and $W_{J''}$ with (part of) its associated rooted tree $J = (J', J'')$. Right: Part of the (unrooted) tree associated to an unpaired intersection $p \in W_J \cap W_K$.

The notation W_J indicates that J is the rooted tree associated to the Whitney disk W, and we identify rooted labeled trees with non-associative bracketings of the index set formed by the link-components. So $W_{(J',J'')}$ pairs intersections between $W_{J'}$ and $W_{J''}$, and the rooted tree (J', J'') is formed by identifying the roots of J' and J'' to a single vertex and sprouting the rooted edge of (J', J'') from this vertex (Fig. 2, left). The *order* of W_J is defined to be the order of J, which is the number of trivalent vertices in J. In particular, for a singleton index j, the order of W_j is zero and its rooted tree is the edge — j.

Now each tree in $t(\mathcal{W})$ corresponds to one of two kinds of *problems* in $\mathcal{W}$: A transverse intersection $p \in W_J \pitchfork W_K$ is a problem if p is not paired by a higher-order Whitney disk. To such a p is associated the labeled tree obtained by identifying the roots of J and K to a single (interior) point, (Fig. 2, right; and Fig. 3, right). The other kind of problem that can occur is that a Whitney disk W_J in $\mathcal{W}$ may be *twisted*, i.e. have a nonzero relative Euler number $\omega(W_J) \in \mathbb{Z}$ with respect to

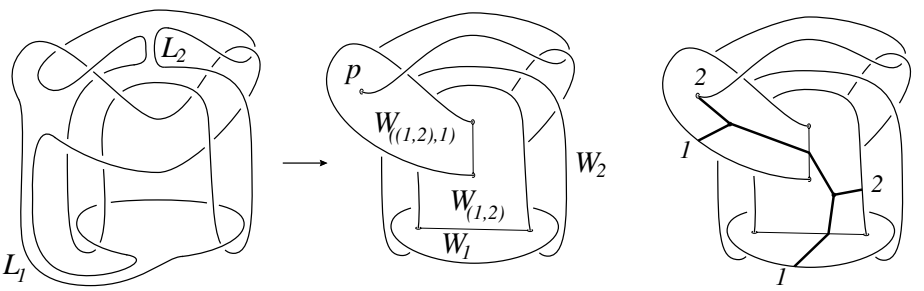

Fig. 3. Left: An internal band sum $L = (L_1, L_2) \subset S^3$ of a Bing-double of the Hopf link. Center: Moving into B^4, the bands have "dissolved" into 1-handles, leaving two 0-handles for each disk W_j bounding L_j, with one 0-handle for W_1 visible as a horizontal disk. The Whitney disk $W_{(1,2)}$ pairs $W_1 \cap W_2$, and there is a single unpaired intersection p between W_2 and the Whitney disk $W_{((1,2),1)}$ pairing $W_{(1,2)} \cap W_1$. Right: The tree associated to p is $t_2 = {}^{\frac{1}{2}}\!\!>\!\!\!-\!\!\!<{}^{\frac{1}{2}}$ from Fig. 1.

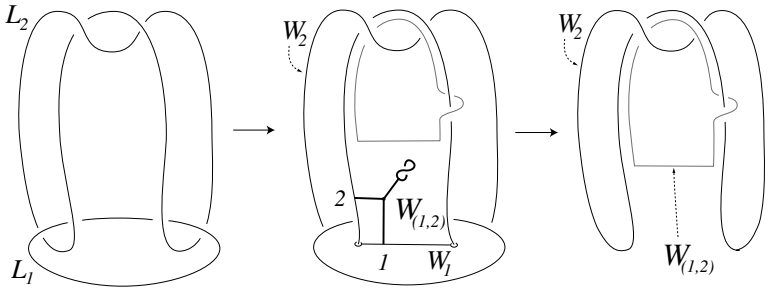

Fig. 4. Left: The Whitehead link $L = (L_1, L_2) \subset S^3$. Center: Moving into B^4, (most of) an embedded twisted Whitney disk $W_{(1,2)}$ pairing $W_1 \cap W_2$, and the associated twisted tree. Right: Continuing into B^4, the rest of $W_{(1,2)}$ and W_2 are described by capping off this unlink with two embedded disks (not shown). The twisting $\omega(W_{(1,2)}) = +1$ is evident in the right-handed twist (see e.g. [5, Sec. 6]).

the standard framing of its boundary (e.g. [5, Sec. 2.2]). This integer $\omega(W_J)$ is called the *twisting* of W_J, and if $\omega(W_J) \neq 0$, then a Whitney move guided by W_J will create new self-intersections, even if W_J happens to be embedded. This second problem contributes to $t(\mathcal{W})$ a tree J^{∞} defined by labeling the root vertex of J with the symbol ∞ standing for "twist".

A Whitney disk with twisting equal to zero is said to be *framed*. Since it can always be arranged (by *splitting* $\mathcal{W}$ if needed [6, Sec. 2.6]) that unpaired intersections p occur between framed disks, we frequently refer to the tree associated to such a p as a *framed* tree, to differentiate from the trees J^{∞} associated to twisted Whitney disks, which we usually call *twisted trees*, or ∞-trees.

Definition 1.3. For a twisted Whitney tower $\mathcal{W}$, define the *intersection forest* $t(\mathcal{W})$ as the disjoint union of (isomorphism classes of) such trivalent, labeled trees, one for each problem in $\mathcal{W}$.

Note that if $\mathcal{W}$ has no problems, i.e. $t(\mathcal{W})$ is empty, then one can do a sequence of (embedded, framed) Whitney moves to produce slice disks for the link L on the boundary.

Remark 1.4. Comparing this definition with the definition of $t(\mathcal{W})$ given in [6, Sec. 2.5], the reader will notice that there the framed trees carry a vertex-orientation and a coefficient ± 1, coming from orientations induced by the original link components (involving *antisymmetry relations*), and the twisted trees carry the integer coefficients coming from the twisting of the corresponding Whitney disks (and independent of disk orientations). We do not need this extra data to define Cochran towers, although it is used in the obstruction theory discussed below.

Remark 1.5. As per Definition 1.6 below, twisted Whitney disks are allowed in *twisted Whitney towers* (e.g. [5, 6]), while earlier papers defined *Whitney towers* with the requirement that Whitney disks be *framed*. For conversational ease we

may sometimes let "Whitney tower" refer to either twisted or framed Whitney towers in a general discussion that applies to both settings, hopefully when no confusion will result.

1.2. *Obstruction theory for twisted Whitney towers*

Definition 1.6. Let $\mathcal{W}$ be a twisted Whitney tower.

(1) The *order* of any tree in $t(\mathcal{W})$ is the number of trivalent vertices. (This applies to both framed and twisted trees.)
(2) If all framed trees in $t(\mathcal{W})$ are of order $\geq n$ and all twisted trees in $t(\mathcal{W})$ are of order $\geq \frac{n}{2}$ then $\mathcal{W}$ is a *twisted Whitney tower of order n*.

This notion of order for twisted Whitney towers arose in our work on Milnor invariants [6], where we discovered that for any order n twisted $\mathcal{W}$ bounded by L the order n framed trees and order $\frac{n}{2}$ twisted trees (for n even) in $t(\mathcal{W})$ contribute to $\mu_{i_1 i_2 \ldots i_{n+2}}(L)$ (with the twisted trees corresponding to certain multi-indices with order 2 symmetry).

More precisely, we showed in [6] how the first non-vanishing μ-invariants can be computed from $t(\mathcal{W})$ modulo certain relations, all of which can be realized via geometric maneuvers preserving the order of $\mathcal{W}$ (without changing its boundary L). Most prominently, the *geometric IHX-relations*, or 4-dimensional Jacobi-identities, can be used to change $t(\mathcal{W})$ by replacing a tree containing an I-shaped subtree with two trees of the same order that only differ locally by H- and X-shaped subtrees, plus a number of trees of higher order [5]. For instance, at the cost of creating higher-order trees, the geometric IHX-relations can be used to modify an order n twisted Whitney tower so that all framed trees in $t(\mathcal{W})$ with two 2-labels and n 1-labels are isomorphic to t_n in Fig. 1, and all twisted trees with one 2-label and $\frac{n}{2}$ 1-labels are isomorphic to $t^{\infty}_{\frac{n}{2}}$ if n is even.

It follows from [6] that nontrivial order n Milnor invariants of L are obstructions to L bounding twisted Whitney towers of order greater than n, and it is not clear how order k Milnor invariants for $k > n$ might be related to intersection forests of order n Whitney towers. Since the specific Milnor invariants we are focusing on in this paper occur in multiple orders and can be lifted to the well-defined Cochran invariants $\beta^i(L)$, we arrive at the question of whether we can compute these in terms of twisted Whitney towers bounded by L. The answer, given by Theorem 1.1, is that indeed we can, and we simply add the twistings of certain Whitney disks. However, there is a catch: For a Cochran tower $\mathcal{W}$, the types of trees in $t(\mathcal{W})$ must be restricted in new ways!

1.3. *Cochran towers*

To define Cochran towers, we start with the following definitions for Whitney towers bounded by 2-component links, so all trees are labeled by 1, 2 or ∞.

Definition 1.7. A trivalent labeled tree is called β-bad if at most one univalent vertex is not labeled by 1, or if it isomorphic to t_i in Fig. 1 for some i.

Definition 1.8. A twisted Whitney tower $\mathcal{W}$ is a *Cochran tower of order n* if all framed β-bad trees in $t(\mathcal{W})$ are of order $> n$ and all twisted β-bad trees in $t(\mathcal{W})$ are of order $> \frac{n}{2}$. Consequently, a twisted Whitney tower $\mathcal{W}$ is a *Cochran tower of infinite order* if $t(\mathcal{W})$ does not contain any β-bad trees at all.

We observe that:

- The framed β-bad trees are the framed trees having no 2-labels, the framed trees having a single 2-label, and the t_i.
- The twisted β-bad trees are the twisted trees having no 2-labels.

Illustrations are given in Examples 1.11, 1.12 and 2.1 showing how the presence of lower-order β-bad trees in a twisted Whitney tower can create indeterminacies in the computation of higher-order β-invariants.

The order zero trees $1 \text{—} 2$ in $t(\mathcal{W})$ give the linking number of L, while $1 \text{—} \overset{1}{\prec}_{1}$ and $1 \text{—} \overset{1}{\prec}_{\infty}$ are responsible for $\mathrm{Arf}(L_1)$. This explains the easy direction in the first sentence of Theorem 1.1. The hard direction is proven by applying all the tricks that we learned over the years, on how to raise the order of a Whitney tower. In particular, those trees that vanish modulo IHX can be exchanged for intersections of arbitrarily high order without changing the β^i.

Remark 1.9. Notice that, via Theorem 1.1, a Cochran tower of order $2k$ allows the computation of $\beta^i(L)$ for all $i \leq k$, i.e. simultaneous extraction of invariants from a Whitney tower at multiple orders, in contrast with the order n Milnor invariants which require L to bound an order n twisted Whitney tower. To our knowledge this is the first example of this kind of computation. Moreover, in an infinite order Cochran tower, the β^i can be computed in all orders, as in the next Example 1.10.

Example 1.10 (The Whitehead link). Figure 4 describes a twisted Whitney tower $\mathcal{W}$ bounded by a (positive) Whitehead Link L, with $t(\mathcal{W}) = 1 \text{—} \overset{2}{\prec}_{\infty} = t_1^{\infty}$. Since $t(\mathcal{W})$ contains no β-bad trees, $\mathcal{W}$ is a Cochran tower of infinite order, and it follows from Theorem 1.1 that $\beta^1(L) = 1$, and $\beta^i(L) = 0$ for all $i > 1$. In fact, if one creates Seifert surfaces by adding tubes to the obvious disks bounded by the components of L (in order to make the Seifert surfaces disjoint from the other component), then the derived link $D(L)$ is the same unlink as in the right-hand side of Fig. 4.

Example 1.11 (Infinitely many nontrivial Cochran invariants, and why the linear framed trees with a single 2-label at one end (Fig. 6, lower right) are β-bad). Let L be the link where the first component is a trefoil knot considered as the boundary of a genus one Seifert surface with a $+1$ twist in each band. The second component is a meridian to one of the two bands. See Fig. 5,

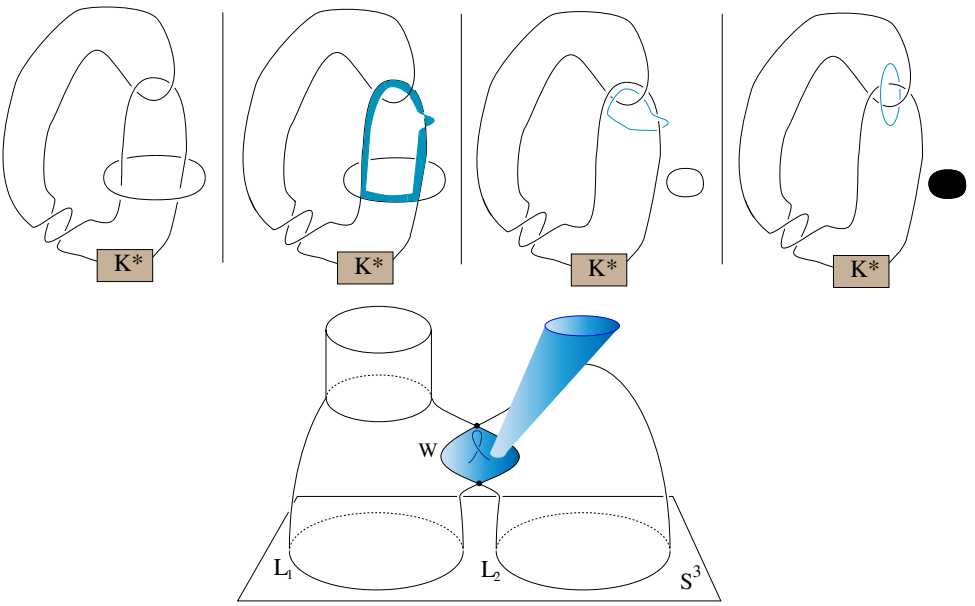

Fig. 5. One stage in the construction of a Cochran tower for the link in Example 1.11.

and [2, Example 4.6]. The derived link $D(L)$ is actually the same link but with the second component linking the other band, and this is isotopic to the original link. So $D(L) = L$. Similarly to the Whitehead link, the Seifert surfaces for L intersect in a knot with twisting $+1$, so $\beta^i(L) = 1$ for all i.

Now we build a Cochran tower of arbitrarily high order for the link. First tie the mirror image of the trefoil (denoted K^* in Fig. 5) inside a small ball in the first component. This will not alter the Cochran invariants. Pulling the second component off the first component we get two intersections which are paired by a $+1$-twisted Whitney "annulus". The free boundary of this annulus is a meridian to the second band and hence the resulting link is isotopic to the original. This is schematically pictured at the bottom of Fig. 5.

One can then iterate this construction arbitrarily, concatenating this basic homotopy repeatedly, and creating a single $+1$-twisted Whitney disk whose associated tree is t_i^{∞} at each ith iteration. To get a twisted Whitney tower one eventually needs to *not* pair the two intersection points, see Fig. 6 for the case of three iterations. The first component can then be capped off with a ribbon disk (because we added the inverse K^*). After k iterations the resulting intersection forest $t(\mathcal{W})$ will contain single ∞-trees t_i^{∞} associated to the $+1$-twisted Whitney disks for each $i \leq k$, plus two order k framed trees associated to the final two unpaired points at the top of the tower. These framed trees are β-bad, since they have a single 2-label and $(k+1)$ 1-labels, so we can only compute Cochran invariants up to a finite order

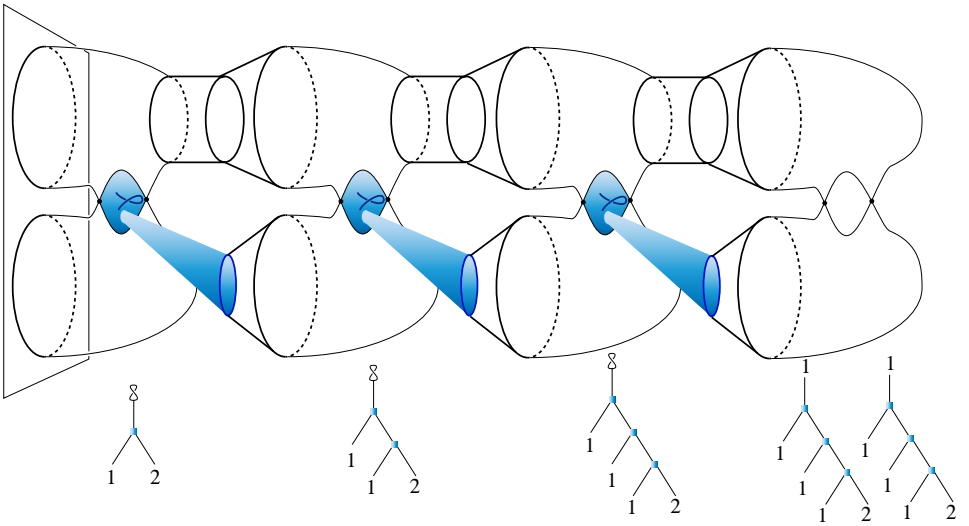

Fig. 6. Iterating the construction to get a Cochran tower of arbitrary order.

using this tower. Iterating the construction $2k+1$ times would yield an order $2k$ Cochran tower C, with the t_i^{∞} trees in C demonstrating that $\beta^i(L) = 1$, for $i \leq k$.

This also illustrates why the linear framed trees at the last step (Fig. 6, lower right) must be considered β-bad.

Example 1.12 (The twisted trees $(1,1,\ldots,1)^{\infty}$ are β-bad). We start with two illustrations of why $(1,1)^{\infty} = 1\!\!-\!\!\prec^{1}_{\infty}$ is β-bad. The disjoint union of a trefoil knot and an unknot bounds a twisted Whitney tower with the single intersection tree $1\!\!-\!\!\prec^{1}_{\infty}$, and the Cochran invariants of this link are trivial. On the other hand, we claim that there is a link with a Whitney tower that has this same single intersection tree, but which has infinitely many nontrivial Cochran invariants. Thus the presence of a tree of this form contributes indeterminacies to the tower which affect the Cochran invariants, and therefore $1\!\!-\!\!\prec^{1}_{\infty}$ must be considered β-bad. The link L is depicted in Fig. 7.

The intersection knot (shown in green in Fig. 7) is untwisted, so $\beta^1(L) = 0$; but $\beta^i(L) = 1$ for all $i \geq 2$, since $D(L)$ is similar to the link with infinitely many Cochran invariants from Example 1.11: The only difference is that the second component of $D(L)$ is a meridian to the Seifert surface of a figure 8 knot instead of a trefoil, but the discussion of Example 1.11 applies to this case, too.

A twisted Whitney tower $\mathcal{W}$ bounded by L with $t(\mathcal{W})$ consisting of a single β-bad tree $1\!\!-\!\!\prec^{1}_{\infty}$ can be constructed as the trace of the null-homotopy of L into B^4 which pulls apart the clasps of L_1, creating a pair of self-intersections in W_1

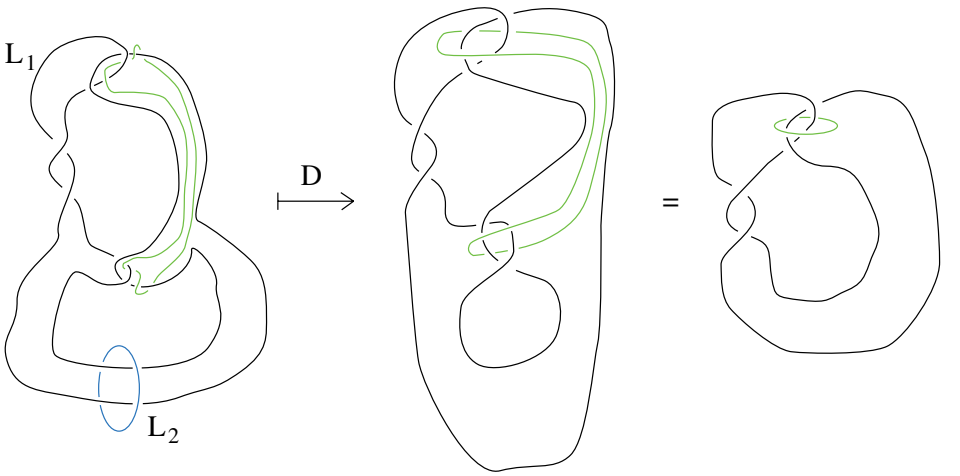

Fig. 7. A link $L = (L_1, L_2)$ with $\beta^i(L) \neq 0$ for all $i \geq 2$. L bounds a twisted Whitney tower $\mathcal{W}$ with $t(\mathcal{W})$ consisting of a single β-bad tree $1\!-\!\!\prec^{1}_{\infty}$.

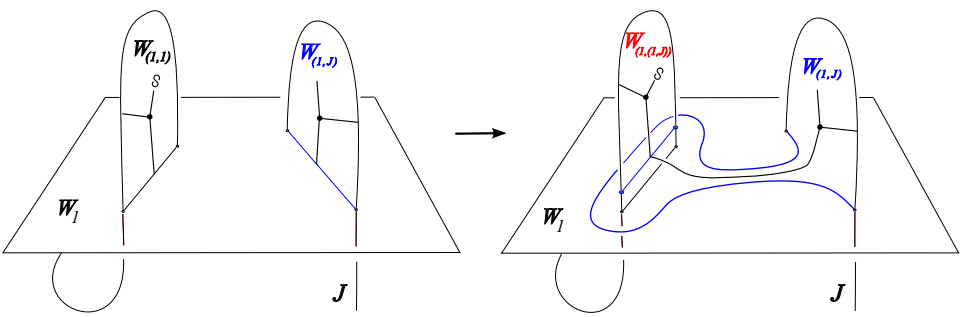

Fig. 8. Left: A twisted Whitney disk $W_{(1,1)}$, and a clean Whitney disk $W_{(1,J)}$ (which can be created for any J by finger moves). Right: Pushing a collar of $W_{(1,J)}$ over $W_{(1,1)}$ gives rise to a new twisted Whitney disk $W_{(1,(1,J))}$ which could create twisted trees t_i^∞ for appropriate J. For instance, t_2^∞ could be created in the case $J = 2$, and t_3^∞ could be created in the case $J = (1, 2)$.

which admit an embedded 1-twisted Whitney disk $W_{(1,1)}$. The interior of $W_{(1,1)}$, and the interior of W_2 bounded by L_2 are both free of intersections.

Figure 8 gives a general illustration how the presence of a twisted $W_{(1,1)}$ contributing $1\!-\!\!\prec^{1}_{\infty}$ to $t(\mathcal{W})$ allows the creation of an arbitrary number of t_i^∞-trees by manipulating $\mathcal{W}$.

The same construction as in Fig. 8, but with a twisted $W_{(1,1,\ldots,1)}$ in the place of $W_{(1,1)}$ (and with $J = 2$), shows how the presence of $(1, 1, \ldots, 1)^\infty \in t(\mathcal{W})$ also allows the creation of an arbitrary number of t_i^∞-trees.

2. Computing Cochran's β^i via Work of Kojima and Kirk

After describing the relationship between the Cochran invariants and an invariant of Kojima, this section applies related computational techniques used by Kirk to give a family of examples which further illustrate how the presence of β-bad trees in $t(\mathcal{W})$ can lead to indeterminacies in the computation of the β^i from $\mathcal{W}$. This section is not used in the proofs of our main results, which are given in Sec. 3.

Cochran shows in [2] that his β^i invariants are related to an invariant due to Kojima, called the η-invariant [16]. Consider the infinite cyclic cover of the complement of a knot L_1, denoted Y, and let $\Delta(t)$ be the symmetrized Alexander polynomial of L_1. For a link $L = (L_1, L_2)$ with linking number zero, let z be a lift of L_2 to Y, z_0 a nearby lift of an untwisted parallel of L_2 and t_* a generator of the covering transformation group.

Then $\Delta(t_*)$ kills $z \in H_1(Y)$, so $\Delta(t_*)(z) = \partial d$ for some 2-chain d in Y and Kojima's η function can be defined as

$$\eta(L) = \sum_{n=-\infty}^{\infty} \frac{1}{\Delta(t)}(z_0 \cdot t_*^n d)t^n \in \mathbb{Z}[t, t^{-1}].$$

Cochran proves that under the change of variables $x = (1-t)(1-t^{-1})$, one has

$$\eta(L) = \sum_{i=1}^{\infty} \beta^i(L)x^i \in \mathbb{Z}[|x|].$$

An important consequence of this formula (with the Alexander polynomial in the denominator) is the existence of many links L (with $\Delta(t) \neq 1$) for which $\beta^i(L) \neq 0$ for *infinitely many* i.

In certain cases this can be related to Wall's self-intersection invariant via a procedure for constructing disjointly immersed surfaces from links used by Kirk to study link maps in the 4-sphere [15]: Consider the special case where L_1 is unknotted, and think of the link L as lying in the upper boundary of $S^3 \times I$. Because the linking number is zero, there is a homotopy of second component L_2, in the complement of L_1, that ends with the trivial link. Thinking of the track of this homotopy as an immersed annulus A in the complement of $L_1 \times I$, one can consider the equivariant intersection number of A with a parallel copy, $\lambda(A, A) \in \mathbb{Z}[t, t^{-1}]$. Both copies of A are immersed annuli in the 4-manifold $(S^3 \smallsetminus L_1) \times I \simeq S^1$ with fundamental group $\mathbb{Z}$, generated by t.

Note that setting $t = 1$ gives zero because one gets the intersection number of an annulus (rel. boundary) in a manifold with no relative second homology. Moreover, the right-hand side comes from a Hermitian form and is hence invariant under the involution $a \mapsto \overline{a}$ on the group ring $\mathbb{Z}[t, t^{-1}]$ determined by $\overline{t} := t^{-1}$. It follows that $\lambda(A, A)$ is a polynomial in $x = (1-t)(1-t^{-1})$ and in fact Kirk showed in [15] that

Kojima's η function is given by

$$\eta(L) = \lambda(A, A) \in \mathbb{Z}[t, t^{-1}].$$

This is the usual relation, in the equivariant setting, between the intersection form of a 4-manifold and its linking form on the boundary.

To compute $\lambda(A, A)$, recall Wall's formula for the relation of the *self-intersection* invariant $\mu(A)$ with the intersection with a parallel copy. If the parallel copy has linking numbers zero on the boundaries then we get

$$\lambda(A, A) = \mu(A) + \overline{\mu(A)}.$$

As a consequence, the easiest way to compute the invariants $\beta^i(L)$ for a link L with linking number zero and unknotted component L_1 is as follows: Do crossing changes on L_2 to separate it from L_1, and for each such crossing change record $\pm t^n$, where the sign is determined by the right-hand rule for the crossing. Moreover, n is the linking number with L_1 of the accessory circle of the crossing (which leaves on one sheet of L_2 and returns on the other).

Note that there are two choices for this accessory circle that differ by L_2 itself. Since it links L_1 trivially, this choice is irrelevant. However, the integer n is only defined up to sign since we do not know in which direction we should run the accessory circle. This is the usual indeterminacy of Wall's self-intersection invariant μ, and it disappears when computing the above average over the involution: One has to record $\pm(t^n + t^{-n})$ for each crossing change and sum over all necessary crossing changes.

This leads to a simple computation of

$$\sum_{i=1}^{\infty} \beta^i(L) x^i = \lambda(A, A) = \mu(A) + \overline{\mu(A)} \quad \text{where} \quad x = (1-t)(1-t^{-1}).$$

Note that in this case, only finitely many of the invariants $\beta^i(L)$ can be nontrivial!

We also note that the exact same computation is valid if L_1 is allowed to be knotted but has trivial Alexander polynomial. By Freedman's Theorem [11], L_1 allows a topological concordance in $S^3 \times I$ to the unknot, with complementary fundamental group $\mathbb{Z}$. Therefore, L_2 bounds an immersed annulus A in the complement of the concordance (leading to the trivial link) and we can again compute $\mu(A)$.

The formulas above extend to this setting because 0-surgery on L_1 is the boundary of the complement of (an open neighborhood of) the slice disk in D^4. Moreover, Kojima's function is the boundary value of the $\mathbb{Z}$-equivariant linking form associated to the intersection form on this 4-manifold.

Example 2.1 (The tree t_2 is β-bad). This example uses clasper-surgery to illustrate in detail how the presence of the framed tree $t_2 \in t(\mathcal{W})$ leads to indeterminacies in the computation of higher-order Cochran invariants in several orders at once. As described in Theorem 3.1 (Section 3 and [8]), the tree-types of clasper

surgeries on an unlink correspond to trees in the intersection forest of a Whitney tower bounded by the resulting link.

Consider the clasper (see Sec. 3.2) of type t_2 pictured below on the left:

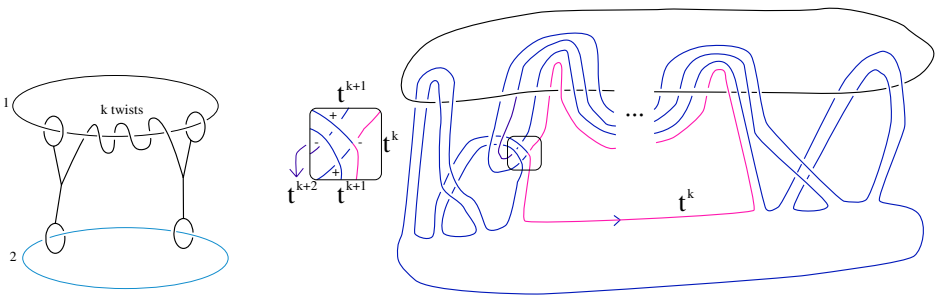

Let L^k be the link obtained from this surgery, pictured on the right. We calculate $\eta(L^k)$ via the Wall self-intersection invariant using the procedure just outlined. There are four crossing changes of component 2, pictured in the small box, which will turn L^k into an unlink. The accessory arc for one of these crossings is colored pink, and one can see that it links with component one k times, and so is assigned $-t^k$, the negative coefficient coming from the fact that the crossing is a negative crossing. One similarly calculates the accessory circles for the other three crossings getting $-t^{k+2}, t^{k+1}, t^{k+1}$. Thus $\eta(L^k) = -(t^k + t^{-k}) + 2(t^{k+1} + t^{-k-1}) - (t^{k+2} + t^{-k-2})$. Converting this to a power series in $x = (1-t)(1-t^{-1})$ we have the following table of polynomials in x:

k	$\sum_{i=1}^{\infty} \beta^i(L^k)x^i$
-1	$2x$
0	$2x - x^2$
1	$2x - 4x^2 + x^3$
2	$2x - 9x^2 + 6x^3 - x^4$
3	$2x - 16x^2 + 20x^3 - 8x^4 + x^5$
4	$2x - 25x^2 + 50x^3 - 35x^4 + 10x^5 - x^6$

The fact that $\beta^1(L^k) = 2$ for all these examples comes from the fact that the link bounds a Whitney tower with the single β-bad tree t_2, and this contributes 2 to the Sato–Levine invariant μ_{1122}. On the other hand, the higher $\beta^i(L^k)$ are not constant on this class of links, so surgery on a t_2-clasper will not have a predictable effect on the Cochran invariants.

3. Proof of Theorem 1.1

After fixing some terminology and notation, this section will prove Theorem 1.1 by combining Habiro's clasper surgery techniques [14] with the Whitney tower obstruction theory [5] and the following decomposition theorem from [8]:

Theorem 3.1. *If a link L bounds a (twisted) Whitney tower $\mathcal{W}$ then there is a finite sequence of concordances and simple (twisted) clasper surgeries from the unlink to L such that the tree-types of the clasper surgeries correspond to the trees in $t(\mathcal{W})$, with the twistings on twisted claspers corresponding to the twistings on the twisted Whitney disks of $\mathcal{W}$.* □

A more precise description of the relationship between clasper concordance and Whitney towers is given in [8].

See Sec. 3.2 below for clasper terminology and conventions, including tree-types and twisted claspers.

3.1. *Outline of proof of Theorem* 1.1

The proof of the second statement of Theorem 1.1, that Cochran invariants up to order k can be calculated from the intersection forest of a Cochran tower of order $2k$, will proceed as follows: By Theorem 3.1, the Cochran tower implies the existence of a sequence of clasper surgeries and concordances from the unlink U to the link L, where the tree-types and twistings of the claspers are the same as those in the intersection forest of the Cochran tower. Since the Cochran invariants are concordance invariants, we then observe how they change under clasper surgeries: On the one hand, a t_i^{∞}-surgery will be shown to leave β^j unchanged for $j \neq i$, and to change β^i by the twisting $\omega(\Gamma)$ of the twisted t_i^{∞}-clasper Γ in Corollary 3.15. On the other hand, clasper surgeries corresponding to all other trees allowed in an order $2k$ Cochran tower will be shown to preserve the β^i for all $i \leq k$ in Corollary 3.10 and Propositions 3.17, 3.16, and 3.18. So $\beta^k(L) - \beta^k(U)$ will be the weighted count of t_i^{∞} trees as claimed.

The proof of the first statement of Theorem 1.1, that a link L bounds a Cochran tower of order $2k$ in B^4 if and only if L has vanishing linking number and the Arf invariant of L_1 vanishes, will follow from the fact that the β-bad trees in the intersection forest of a Whitney tower bounded by L can all be eliminated by twisting constructions or geometric IHX-constructions, at the cost of only creating higher-order trees [5]. See Sec. 3.5.

3.2. *Clasper conventions*

We will be using Habiro's *clasper surgery* techniques, for details see [9, 14]. We adopt the following terminology from [9], and essentially the same notion of twisted claspers from [7]. Although claspers are surfaces, we follow the customary identification of a clasper with its 1-dimensional spine, which is a framed unitrivalent graph.

All of our claspers will be *tree claspers*. A clasper $\Gamma \subset S^3$ is *capped* if all the leaves of Γ bound disjointly embedded disks (the *caps*) into $S^3 \backslash \Gamma$ (so the leaves of a capped clasper are unknotted). A cap for a clasper on a link is called a *simple cap* if it is 0-framed and intersects the link in a single point. A *simple tree clasper* is a

capped tree clasper on a link such that each cap is simple. A *twisted tree clasper* is a capped tree clasper on a link such that all caps are simple except for one k-framed cap, for some integer $k \neq 0$, whose interior is disjoint from the link. If Γ is a twisted tree clasper, this *twisting number* k of Γ will be denoted $\omega(\Gamma)$.

A leaf is *simple* if it bounds a simple cap and *twisted* (k*-twisted*) if it bounds a twisted cap. If a clasper leaf bounds a 0-framed cap which is disjoint from all link components, then the leaf and the cap are said to be *clean*.

By the *interior* of a clasper, we will mean the complement of the leaves (and caps), which in our case is a unitrivalent tree. For simple and twisted tree claspers this tree is labeled exactly like the intersection trees in $t(\mathcal{W})$: Univalent vertices are labeled by the components that the corresponding caps intersect, or by the ∞ symbol if that leaf is twisted.

A clasper with tree type T will be called a T*-clasper*. In particular, if T is a twisted tree, then a T-clasper is a twisted clasper.

Surgery with a specified root: For Γ a T-clasper in the complement of a link L, we denote by L_Γ the link obtained by surgery on Γ. This surgery is really a diffeomorphism of S^3 which can change the link's isotopy class. The diffeomorphism and concomitant modification of L can be realized in many different ways. For example, if one designates a univalent vertex of T as the root, one can realize L_Γ as a modification of L which takes place in a neighborhood of Γ and the root cap; the other caps are not involved. More precisely, only strands of L which intersect the root cap will be modified, and this modification will be supported near Γ. A surgery on a T-clasper will be called a T*-surgery*.

The Zip construction: Suppose a leaf of a clasper Γ bounds a cap Δ which may intersect the link L. The *zip construction* [9, 14] cuts the cap Δ into two pieces, such that surgery on Γ is equivalent to surgery on the union of two daughter claspers, both of the same tree type as Γ, each of which inherits one piece of Δ as a cap, while the other leaves are parallel copies of the originals. The two daughter claspers are embedded in a neighborhood of Γ; their edges are disjoint from any caps that each may have.

3.3. *The Cochran β^i-invariants*

We focus on the setting of classical 2-component links modulo concordance, and refer the reader to [2, 3] for details on the full generality of the Cochran β^i-invariants.

Definition 3.2 ([2]). Let $L = (L_1, L_2)$ be a 2-component link with zero linking number, and let Σ_1 and Σ_2 be Seifert surfaces for L_1 and L_2 such that $\Sigma_1 \cap \Sigma_2 = \chi$ is connected (and nonempty). Then the *derived link* $D(L)$ is defined to be the 2-component link (L_1, χ) obtained by replacing L_2 with χ. The knot χ is called the *characteristic curve* of L (and Σ_1 and Σ_2).

It is easy to see that a derived link $D(L)$ also has zero linking number, giving rise to the following infinite sequence of invariants, which Cochran showed to be concordance invariants of L:

Definition 3.3 (*Cochran invariants* [2]). The first Cochran invariant $\beta^1(L) \in \mathbb{Z}$ is defined to be the classical Sato–Levine invariant of L. Inductively, $\beta^i(L)$ is defined to be $\beta^{i-1}(D(L))$ for $i > 1$.

The classical Sato–Levine invariant [19] is equal to the *twisting* $\omega(\chi) \in \mathbb{Z}$ of the framing of the characteristic curve $\chi = \Sigma_1 \cap \Sigma_2$ induced by the normal framing of either surface Σ_j relative to the 0-framing of $\chi \subset S^3$.

Remark 3.4. As shown by Cochran [2], for any $L = (L_1, L_2)$ with trivial linking number, there exist Seifert surfaces Σ_1 and Σ_2 that intersect in a connected closed curve. In the sequel, it will often happen that modification of such a link will temporarily create new components of intersection between the modified Seifert surfaces, and we will use the following procedure for making the intersection connected in a well-controlled way: The modification of L will lead to Σ_1 being extended by boundary-summing into some new genus one surfaces, and the corresponding new intersection curves with the second Seifert surface will be at most a single non-separating circle on each of these new genus one subsurfaces. Choose arcs in the new Σ_1 connecting each new intersection circle to the original characteristic curve $\chi = \Sigma_1 \cap \Sigma_2$. Surgering the second Seifert surface along these arcs has the effect of band-summing the new intersection curves with χ, making the intersection between the new Seifert surfaces connected. Since each of the new circles was non-separating, these arcs can be chosen in such a way that the surgered surface is still orientable. (For example, see Figs. 10–12, where new genus is added to Σ_1 along with new intersection components supported in the new genus.)

3.4. *Claspers and Cochran invariants*

It turns out that the Cochran invariants are *finite type* in a certain sense, ultimately deriving from the fact that the twisting $\omega(K)$ of a framed knot K is a type 1 framed knot invariant, a notion we now define.

Recall that a nowhere-vanishing normal vector field of a smooth knot $K \subset S^3$ is referred to as a *framing* of K. Identifying the image of the vector field with a knot parallel to K, the twisting $\omega(K)$ is the linking number of K with its parallel copy. Here the notation K is abused to describe the *framed knot*.

Definition 3.5 ([1]). A *singular framed knot* (in S^3) is a framed knot, modified to allow finitely many double points or non-tangential zeros of the framing. Any framed knot invariant v can be extended to singular framed knots, by the usual

Vassiliev skein relation for double points and by

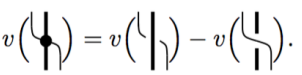

for the non-tangential zeros. A framed knot invariant is said to be *type* k if it vanishes on singular framed knots with $k+1$ double points and non-tangential zeroes.

In other words, for any collection S of $k+1$ disjointly supported crossing changes and local framing changes, a type k framed knot invariant v will satisfy

$$\sum_{S' \subset S} (-1)^{|S'|} v(K_{S'}) = 0$$

where $K_{S'}$ the result of changing K by exactly the elements in S', $|S'|$ is the cardinality of S', and the sum is over all subsets of S. Crossing changes and local framing changes are both examples of homotopies of K supported in balls (where we consider the local framing change to be a crossing change between K and the parallel knot corresponding to the framing), and all homotopies and framing changes can be represented as a sequence of these two basic moves (together with isotopy.) As observed by Goussarov [13], an alternating sum over a collection S of groups of crossing and framing changes (all disjointly supported) can be written as a linear combination over alternating sums of crossing and framing changes. Thus, we have that an invariant is of type k if and only if it vanishes on alternating sums of $k+1$ disjointly supported homotopies.

It is not difficult to show that the twisting $\omega(K)$ is a framed knot invariant of type 1. This has implications for the Cochran invariants since we ultimately compute them using the twisting of the characteristic curve. However, the β^i are not quite type 1 invariants, but rather are only type 1 invariants if one only allows crossing changes on component 2. This leads to the following definition:

Definition 3.6. A link invariant f is *type k with respect to a given component* if, when given any set S of any $k+1$ disjointly supported crossing changes of this component with itself, we have

$$\sum_{S' \subset S} (-1)^{|S'|} f(L_{S'}) = 0,$$

where $L_{S'}$ denotes the result of changing L by exactly the homotopies in S', $|S'|$ is the cardinality of S', and the sum is over all subsets of S.

Remark 3.7. As above, we can replace crossing changes with arbitrary homotopies of the given component supported in disjoint balls that avoid the other components.

Remark 3.8. In the above paragraphs, we introduced two distinct notions of finite type invariants, one for framed knots and one for links with a specified component.

There is an obvious way to combine those notions, namely for links with a *specified framed* component. Cochran's derived link is exactly of that type and Proposition 3.9 below holds for all $i \geq 0$ if we define $\beta^0(L) := \omega(L_2)$ for 2-component links with a framing on the specified component L_2. Then the proof below is an induction that naturally extends all the way down to $i = 0$.

Convention: For the rest of the paper we assume that L designates a link for which the β^i-invariants are defined, and that L comes equipped with Seifert surfaces which intersect in a single circle.

Proposition 3.9. *For all $i \geq 1$, the Cochran β^i-invariants of $L = (L_1, L_2)$ are type 1 with respect to L_2.*

Proof. We start by showing that the first Cochran invariant β^1 vanishes on an alternating sum over a pair of disjointly supported crossing changes on L_2. These crossing changes can be realized as a pair of edge-clasper surgeries Γ_j, $j = 1, 2$, where Γ_j has two leaves, each of which is a meridian to L_2 and bounds a cap that intersects Σ_2 in a clasp singularity. We may further assume these caps are disjoint from Σ_1.

Now the single edge of each Γ_j will crash through sheets of both Seifert surfaces (Fig. 9, left). Pushing these sheets off the edge of Γ_j will create ribbon singularities with Σ_2. When Σ_2 is pushed off, these new singularities can be resolved using the standard ribbon singularity resolution, increasing the genus of Σ_2 (Fig. 9, second from left, blue).

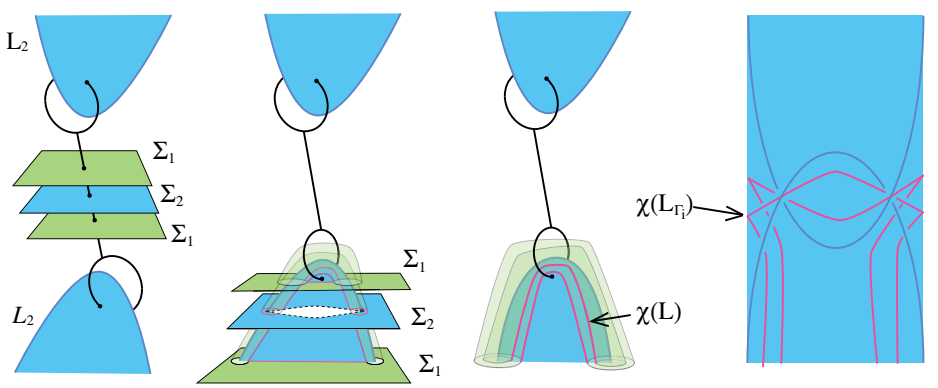

Fig. 9. (Color online) Left: An edge-clasper Γ_j on L_2 which has two interior intersections with Σ_1 and one interior intersection with Σ_2. Second from left: Pushing the surface sheets off the edge creates more genus and intersections among the surfaces. Third from left: Zooming in near the leaf, nested tubes from Σ_1 containing arcs of the characteristic curve $\chi(L)$ (red) of L also link this leaf. Right: A local picture of the characteristic curve $\chi(L_{\Gamma_j})$ (red) for L_{Γ_j} near the clasp created by Γ_j-surgery.

When Σ_1 is pushed off, pairs of algebraically cancelling intersections between Σ_1 and L_2 will be created. Resolve these by adding tubes in a neighborhood of L_2 (Fig. 9, second from left, green). This will create new circles of intersections between Σ_1 and Σ_2 that travel through one leaf of each Γ_j. Apply the procedure of Remark 3.4 (as needed) to modify the surfaces in a neighborhood of $\Sigma_1 \cup \Sigma_2$ so that $\Sigma_1 \cap \Sigma_2$ becomes connected. At this point we have modified the surfaces so that they are disjoint from the interior of the claspers, but one leaf of each clasper may link the characteristic curve $\chi(L)$ of L. Note that we have pushed the surface sheets across the same leaf in each clasper so the other leaf can be used as a root for surgery.

Now we consider the new characteristic curves $\chi(L_{\Gamma_j})$ for the L_{Γ_j}. Each L_{Γ_j} has an added clasp, and Σ_2 can be locally modified into a Seifert surface for component 2 of L_{Γ_j}, which locally looks like a disk with a twisted band attached, as in the far right-hand side of Fig. 9. The tubes from Σ_1 consist of meridians to L_2 which now travel over this band. So the characteristic curve for L_{Γ_j} is formed by a homotopy H_j of $\chi(L)$ in a neighborhood of Γ_j which will introduce the local twisting and crossing changes illustrated in the right-most picture of Fig. 9. Specifically, for each tube of Σ_1 passing through the leaf of Γ_j the homotopy H_j will create one twist in the characteristic curve framing, and crossing changes with the other arcs in tubes through the same leaf.

Since the twisting ω is a type 1 framed knot invariant, this shows that $\beta^1(L)$ is type 1 with respect to L_2:

$$\begin{aligned}
&\beta^1(L) - \beta^1(L_{\Gamma_1}) - \beta^1(L_{\Gamma_2}) + \beta^1(L_{\Gamma_1\cup\Gamma_2}) \\
&\quad = \omega(\chi(L)) - \omega(\chi(L_{\Gamma_1})) - \omega(\chi(L_{\Gamma_2})) + \omega(\chi(L_{\Gamma_1\cup\Gamma_2})) \\
&\quad = \omega(\chi(L)) - \omega(\chi(L)_{H_1})) - \omega(\chi(L)_{H_2})) + \omega(\chi(L)_{H_1\cup H_2})) = 0.
\end{aligned}$$

Now to see that $\beta^i(L)$ is type1 with respect to L_2 for all $i > 1$, observe that by the above constructions the H_j are disjointly supported homotopies of the second component of $D(L)$. So $D(L_{\Gamma_j}) = D(L)_{H_j}$, and by induction

$$\begin{aligned}
&\beta^i(L) - \beta^i(L_{\Gamma_1}) - \beta^i(L_{\Gamma_2}) + \beta^i(L_{\Gamma_1\cup\Gamma_2}) \\
&\quad = \beta^{i-1}(D(L)) - \beta^{i-1}(D(L)_{H_1}) - \beta^{i-1}(D(L)_{H_2}) + \beta^{i-1}(D(L)_{H_1\cup H_2}) = 0. \quad \square
\end{aligned}$$

Corollary 3.10. *If T is any framed tree having at least three 2-labels, or any twisted tree having at least two 2-labels, and Γ is any T-clasper on L, then $\beta^i(L_\Gamma) = \beta^i(L)$ for all $i \geq 1$.*

Proof. We realize L_Γ by doing surgery, using one of the 2-labeled leaves of Γ as the root. Now L_Γ has strands of component 2 (but not component 1) going through a neighborhood of Γ.

Given a non-root simple leaf ℓ of Γ whose cap Δ intersected L_2, in L_Γ there are now multiple parallel strands of component 2 running through where ℓ used to be.

Let H_ℓ be a homotopy that pushes the strand of component 2 that intersected Δ across the parallel strands. Then $(L_\Gamma)_{H_\ell}$ is actually surgery of L along a clasper that has a clean cap. This implies that $(L_\Gamma)_{H_\ell} = L$.

Given a non-root twisted leaf ℓ^∞ of Γ, let H_{ℓ^∞} be a homotopy that straightens the strands of component 2 traveling through where ℓ^∞ used to be. By similar reasoning as in the previous paragraph $(L_\Gamma)_{H_{\ell^\infty}} = L$.

Moreover, given several of these two types of homotopies supported near different leaves of the clasper, modifying L_Γ by any nonempty subset will yield L, as it is equivalent to making more than one leaf bound a clean cap in the complement of both components. The hypotheses of the corollary imply that we may find two disjointly supported homotopies of this form. Thus, by the Proposition 3.9: $\beta^i(L_\Gamma) - \beta^i(L) - \beta^i(L) + \beta^i(L) = 0$. □

Recall that t_k^∞ is the order k twisted tree having k 1-labels, one 2-label and one ∞-label as in Fig. 1.

Proposition 3.11. *If Γ is a t_1^∞-clasper on L, then $L_\Gamma = B_{\omega(\Gamma)} \#_b L$, where B_ω is the ω-twisted Bing double of the unknot, and $\#_b$ denotes a component-wise band-sum.*

Here the bands guiding the band-sum run along the two edges of Γ whose caps intersect L.

Proof. Here is a picture of surgery on a t_1^∞-clasper with twisting ω.

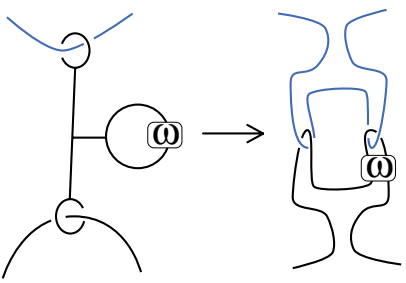

This is exactly a band sum with $B_{\omega(\Gamma)}$ as claimed. □

Corollary 3.12. *Let Γ be a t_1^∞-clasper on L. If $i \geq 2$, then $\beta^i(L) = \beta^i(L_\Gamma)$, and $\beta^1(L_\Gamma) = \omega(\Gamma) + \beta^1(L)$.*

Proof. The invariants β^i are additive under arbitrary band sum [2, Theorem 5.6], and their value on B_ω are 0 for $i \geq 2$ and ω for $i = 1$. □

Lemma 3.13. *Let T be a twisted tree with only one 2-label. Then for any T-clasper Γ on L, there exists a T-clasper Γ' on L satisfying the following:*

(1) $\beta^i(L_\Gamma) = \beta^i(L_{\Gamma'})$ *for all* $i \geq 1$.
(2) $\omega(\Gamma') = \omega(\Gamma)$.

(3) *The interior of* Γ' *is disjoint from possibly new Seifert surfaces for* L*, and taking the derived link* $D(L)$ *with respect to the new Seifert surfaces is the same as the derived link taken with respect to the original Seifert surfaces.*

Proof. Starting with the original Seifert surfaces Σ_1, Σ_2 for L, we may assume that each j-labeled cap of Γ meets Σ_j in a clasp singularity, but is otherwise disjoint from both Seifert surfaces. First, we find a new T-clasper Γ' such that the interior of Γ' is disjoint from Σ_2, and $L_{\Gamma'}$ has the same Cochran invariants as L_Γ. If the interior of Γ is not already disjoint from Σ_2, then do surgery on Γ, using the 2-labeled leaf as the root. Then we can create the following two disjointly supported homotopies of the second component of L_Γ. The homotopy H_1 untwists the twisted strands of component 2 near what was the twisted leaf of Γ. The homotopy H_2 pushes Γ across the boundary of Σ_2, removing all intersections with Σ_2. (Here H_2 is really pushing strands of component 2 that are contained in a neighborhood of where Γ was before the surgery.) Now $\beta^i(L_\Gamma) - \beta^i((L_\Gamma)_{H_1}) - \beta^i((L_\Gamma)_{H_2}) + \beta^i((L_\Gamma)_{H_1 \cup H_2}) = 0$. But $(L_\Gamma)_{H_1} = L = (L_\Gamma)_{H_1 \cup H_2}$, so we get that $\beta^i(L_\Gamma) = \beta^i((L_\Gamma)_{H_2})$. Now $(L_\Gamma)_{H_2} = L_{\Gamma'}$, where Γ' is a T-clasper (the result of applying H_2 to Γ) that has no interior intersections with Σ_2.

So far we have produced a T-clasper Γ' on L satisfying $\forall i\ \beta^i(L_\Gamma) = \beta^i(L_{\Gamma'})$ which does not have interior intersections with Σ_2. If the interior of Γ' is also disjoint from Σ_1 we are done.

Otherwise, suppose that the interior of Γ' is not disjoint from Σ_1. Since only one leaf of Γ' links L_2, we can push all of these interior intersections away from the 2-labeled leaf and toward the 1-labeled leaves and twisted leaf, eventually pushing sheets of Σ_1 off of these leaves and resolving the new self-intersections of Σ_1 via the standard ribbon singularity resolution, as in the proof of Proposition 3.9. This modification of Σ_1 does not affect the original characteristic curve χ. Moreover, L and Γ' are unchanged by this procedure. So $\forall i\ \beta^i(L_\Gamma) = \beta^i(L_{\Gamma'})$, and we have removed all intersections between the interior of Γ' and both Seifert surfaces. □

Proposition 3.14. *Let* Γ *be a* t_k^∞*-clasper on* L *for* $k \geq 2$*. Then there exists a* t_k^∞*-clasper* Γ' *on* L *such that* $\beta^i(L_\Gamma) = \beta^i(L_{\Gamma'})$*, for all* $i \geq 1$*, and a* t_{k-1}^∞*-clasper* Γ'' *on* $D(L)$ *such that* $D(L_{\Gamma'}) = D(L)_{\Gamma''}$ *and* $\omega(\Gamma) = \omega(\Gamma'')$*.*

Proof. Take Γ' to be the clasper guaranteed by Lemma 3.13, with the interior of Γ' disjoint from both Seifert surfaces for L, as in the upper left of Fig. 10. The clasper calculus and constructions illustrated in Fig. 10 show how splitting off a Y-clasper from Γ' yields the desired t_{k-1}^∞-clasper Γ'' on $D(L)$. □

We get the following corollary of Proposition 3.14 and Corollary 3.12 by induction:

Corollary 3.15. *If* Γ *is a* t_k^∞*-clasper on* L*, then* $\beta^k(L_\Gamma) = \omega(\Gamma) + \beta^k(L)$*, and if* $i \neq k$*, then* $\beta^i(L) = \beta^i(L_\Gamma)$*.*

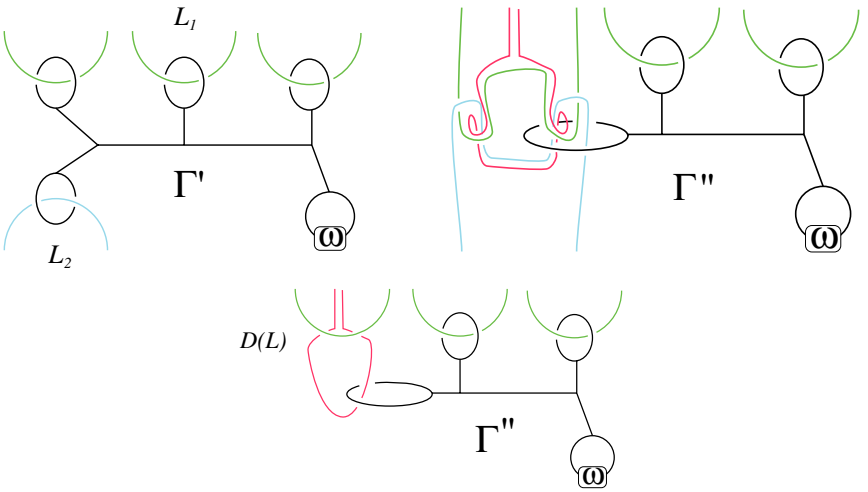

Fig. 10. (Color online) The case $k = 3$ from the proof of Proposition 3.14: Upper left: The interior of the t^∞_3-clasper Γ' on L is disjoint from both Seifert surfaces (which are suppressed from view). Green arcs are from L_1 and the blue arc is from L_2. Upper right: Splitting off a Y-subtree by inserting a Hopf-pair of leaves into an edge of Γ' and then doing the Y-surgery on L yields L_Y, which is isotopic to L, and a further surgery on the clasper Γ'' would also yield $L_{\Gamma'}$. The red curve shows the intersection between extensions of the Seifert surfaces for L to Seifert surfaces for L_Y (which has been band-summed to the original characteristic curve, as per Remark 3.4). Lower: The clasper Γ'' is a t^∞_2-clasper on $D(L)$, with $\omega(\Gamma'') = \omega(\Gamma')$. Removing Γ'', the new characteristic curve is isotopic to the old one since it contracts along the band. Hence $D(L_{\Gamma'}) = D(L)_{\Gamma''}$.

Proposition 3.16. *If T is a twisted tree which has exactly one 2-label, such that $T \neq t^\infty_k$ for any k, and Γ is any T-clasper on L, then $\beta^i(L) = \beta^i(L_\Gamma)$ for all $i \geq 1$.*

Proof. By Lemma 3.13, we may assume that the interior of Γ is disjoint from the surfaces Σ_1 and Σ_2. The proof will proceed by covering the three possible cases:

A pair of vertices is said to be *dual* if they are connected by edges to a common trivalent vertex.

Case 1: The 2-labeled vertex of T is not dual to any univalent vertex. Insert Hopf-pairs of leaves into two edges of Γ to split Γ into three claspers, one of which is a Y-clasper having a leaf linking L_2. (Surgery on all three of these claspers is equivalent to surgery on Γ.) The top pictures in Fig. 11 show the result of performing the Y-surgery using the leaf linking L_2 as a root. As can be seen in the upper right picture of Fig. 11, the original Seifert surfaces for L can be extended to Seifert surfaces on L_Γ without changing the characteristic curve: The indicated (blue) genus one piece is added to Σ_2 in the complement of the other two claspers. And after performing the other two clasper surgeries, Σ_1 (green) can be extended to a Seifert surface for the first component of L_Γ inside a neighborhood of these claspers, so no new intersections are created. Thus $D(L_\Gamma) = D(L)$ and so Γ-surgery preserves all β^i.

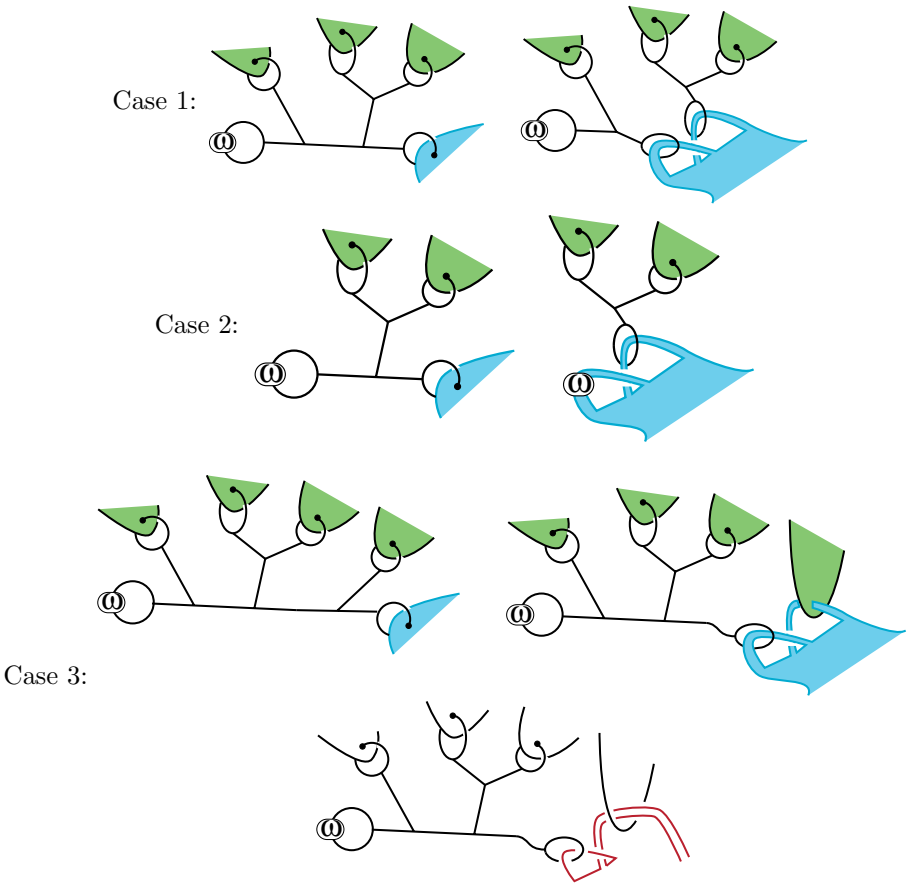

Fig. 11. From the proof of Proposition 3.16.

Case 2: The 2-labeled vertex of T is dual to the ∞-labeled vertex. First note that in this case the 2-labeled vertex cannot also be dual to a 1-labeled vertex, by the assumption $T \neq t_1^{\infty}$. Split Γ into two claspers, one of which is a Y-clasper with one twisted leaf and one leaf linking L_2. The middle pictures in Fig. 11 show the result of performing the Y-surgery using the leaf linking L_2 as a root. As can be seen in the middle right picture of Fig. 11, the original Seifert surfaces for L can be extended to Seifert surfaces on L_Γ without changing the characteristic curve: The indicated (blue) genus one piece with the $\omega(\Gamma)$-twisted band is added to Σ_2 in the complement of the other clasper. And after performing the other clasper surgery, Σ_1 (green) can be extended to a Seifert surface for the first component of L_Γ inside a neighborhood of this clasper, so no new intersections are created. Thus $D(L_\Gamma) = D(L)$ and so Γ-surgery preserves all β^i.

Case 3: The 2-labeled vertex of T is dual to a 1-labeled vertex. Note that in this case the 2-labeled vertex cannot also be dual to the ∞-labeled vertex, by the assumption

$T \neq t_1^{\infty}$. Split Γ into two claspers, Γ' and a Y-clasper with one leaf linking L_2, as in the bottom pictures of Fig. 11. The bottom right picture in Fig. 11 shows the result of performing the Y-surgery using the leaf linking L_2 as a root. A Seifert surface Σ_2' for the new second component can be constructed by adding the indicated (blue) genus one piece to the original Σ_2. Now Σ_2' is in the complement of Γ', but Σ_2' has a ribbon intersection with the original Σ_1. After surgery on Γ' (yielding L_Γ), a Seifert surface Σ_1' for component 1 of L_Γ can be constructed by resolving the ribbon intersection by adding a tube that runs along a subarc of component 2 of L_Γ, and extending Σ_1 in a neighborhood of Γ'. Now the characteristic curve $\chi' = \Sigma_1' \cap \Sigma_2'$ for L_Γ has a new loop which runs over one band of Σ_2' and links Γ'. This new loop is connected to the original χ for L by a band. This new χ' has no additional twists in it, and moreover the derived link $D(L)$ is given by surgery on the clasper Γ' on (L_1, χ') whose tree is of order one less than the order of T. Now the assumption $T \neq t_k^{\infty}$ means that we can proceed inductively, since iterating this reduction will eventually lead to Case 1 or Case 2. □

Recall that t_n is the order n framed tree having two 2-labels and n 1-labels as in Fig. 1.

Proposition 3.17. *If T is a framed tree which has two 2-labels and n 1-labels, for $n \geq 2$, such that $T \neq t_n$, and Γ is any T-clasper on L, then $\beta^i(L) = \beta^i(L_\Gamma)$ for all $i \geq 1$.*

Proof. Consider such a T-clasper on L. Split Γ by inserting Hopf-pairs of leaves into some edges in such a way that yields a union of five claspers $\Gamma_0, \Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4$ with the following properties: Γ_0 has four leaves. A dual pair of leaves of Γ_0 each link leaves of Γ_1 and Γ_2, while the other dual pair of leaves of Γ_0 each link Γ_3 and Γ_4. (A pair of leaves is said to be *dual* if they are connected by edges to a common trivalent vertex.) Moreover, the other leaves of Γ_1 and Γ_2 are meridians to L_1, while there is a single leaf of each of Γ_3 and Γ_4 which is a meridian to L_2. Such a splitting of Γ is possible because $T \neq t_n$.

Now do surgery on $\Gamma_1 \cup \Gamma_2 \cup \Gamma_3 \cup \Gamma_4$, where we take the roots of Γ_3 and Γ_4 to be the leaves linking L_2. This surgery does not change L because each of these claspers has a leaf which bounds a clean cap (although these leaves do link the leaves of Γ_0). Looking at Γ_0 on this link, we see that parallel strands of L_1 go through one dual pair of leaves of Γ_0 whereas parallel strands of L_2 go through the other dual pair.

Now using the zip construction multiple times to reduce the number of strands through the leaves to 1 each, surgery on Γ_0 becomes equivalent to a sequence of surgeries on simple t-claspers, with $t = \overset{1}{1} \!>\!\!-\!\!<\! \overset{2}{2}$. So it suffices to show that such t-surgeries do not affect Cochran invariants.

Consider a t-clasper Γ, with $t = \overset{1}{1} \!>\!\!-\!\!<\! \overset{2}{2}$. We may assume that the caps of Γ only meet the Seifert surfaces in four clasp singularities, and we may assume that the interior of Γ is disjoint from Σ_2 by the following argument: Let L_Γ denote the

result of t-surgery using one of the 2-labeled leaves as the root. Denote by H_1 the homotopy of L_Γ which pushes the strand of component 2 that ran through the non-root leaf across the strands of component 2 that now run along where the non-root leaf used to be. And denote by H_2 the homotopy of component 2 of L_Γ induced by pushing Γ off of Σ_2 (so H_2 moves strands that run along where Γ used to be). Now, applying Proposition 3.9 and noting that $(L_\Gamma)_{H_1} = L = (L_\Gamma)_{H_1 \cup H_2}$ (since the results of H_1 and $H_1 \cup H_2$ are the same as clasper surgeries on L by claspers that each have a clean leaf) gives:

$$0 = \beta^i(L_\Gamma) - \beta^i((L_\Gamma)_{H_1}) - \beta^i((L_\Gamma)_{H_2}) + \beta^i((L_\Gamma)_{H_1 \cup H_2}) = \beta^i(L_\Gamma) - \beta^i((L_\Gamma)_{H_2}).$$

So pushing Γ off of Σ_2 by H_2 gets the interior of Γ disjoint from Σ_2 without affecting the $\beta^i(L)$.

Furthermore, any intersections between Σ_1 and the interior of Γ can all be pushed onto an edge adjacent to a single 2-labeled leaf: Any intersections between Σ_1 and the 1-labeled edges of Γ can be pushed into Σ_1, and the resulting ribbon singularities resolved without affecting the characteristic curve.

One can now draw an explicit picture of the derived link (Fig. 12). The new intersections of Σ_1 and Σ_2 come from the intersections of Σ_1 with the 2-labeled edge of Γ, and are connected by bands to the rest of the characteristic curve. These new curves are all trivial, implying that $D(L_\Gamma) = D(L)$. Furthermore, the twisting of these additions to the characteristic curve are all 0, so that $\beta^1(L) = \beta^1(L_\Gamma)$. These last two facts are sufficient to show that $\beta^i(L) = \beta^i(L_\Gamma)$ for all $i \geq 1$. □

Proposition 3.18. *If T is a framed tree of order $> 2k+1$, or a twisted tree of order $> k$, and Γ is any T-clasper on L, then $\beta^i(L) = \beta^i(L_\Gamma)$ for all $i \geq 1$.*

Proof. First note that for some choice of bands b, the band-sum $L \#_b \bar{L}$ is a slice link, which therefore has vanishing μ-invariants in all orders. In particular all β^i

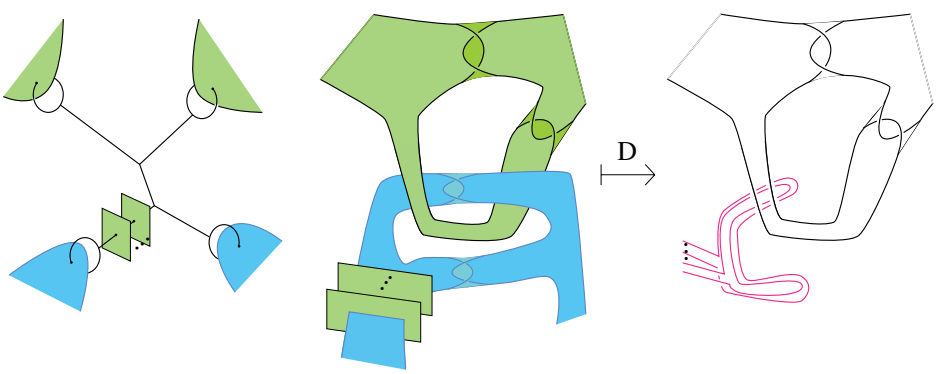

Fig. 12. (Color online) A t-surgery on L, for $t = \frac{1}{1} \gtrless \frac{2}{2}$, with sheets of Σ_1 (green) intersecting the clasper (the blue sheets are from Σ_2). The picture on the right shows $D(L_\Gamma) = D(L)$, with the characteristic curve in red.

vanish [3, Theorem 6.10]. Now consider $(L\#_b\bar{L})_\Gamma = L_\Gamma\#_b\bar{L}$. Since Γ preserves vanishing of μ-invariants up to order $2k$ [14], we have that $0 = \beta^i(L_\Gamma\#_b\bar{L}) = \beta^i(L_\Gamma) - \beta^i(L)$ for $i \leq k$. □

3.5. *Proof of Theorem* 1.1

Finally, we put all the ingredients together to prove Theorem 1.1:

Proof. First, we show that any $L = (L_1, L_2)$ with trivial linking number and $\mathrm{Arf}(L_1) = 0$ bounds a Cochran tower of arbitrarily high order.

Start by taking any pair of properly immersed disks bounding the components of L. Since the linking number is zero, all intersections between the two disks will algebraically cancel, and so can be paired by Whitney disks (which will in general intersect each other and the two immersed disks). By performing local cusp homotopies, it can also be arranged that the disks' self-intersections are all paired by Whitney disks. This yields a twisted Whitney tower $\mathcal{W}$ of order one. If there are no β-bad trees in $t(\mathcal{W})$ then we have a Cochran tower (of infinite order) and we are done. Otherwise, there exist β-bad trees.

Throughout the following arguments for exchanging β-bad trees for higher-order trees, we use the "order-raising" obstruction theory described in [5]. We refer the reader to that paper for details, including orientation and sign conventions giving rise to coefficients for the trees in $t(\mathcal{W})$, as well as descriptions of the various twisting and IHX modifications of $\mathcal{W}$. In particular, whenever a pair of intersections points or twisted Whitney disks contribute isomorphic but oppositely-signed trees to $t(\mathcal{W})$, then that pair can be eliminated at the cost of only creating new higher-order trees in $t(\mathcal{W})$. Discussion of signs will be suppressed in the following constructions.

Any β-bad tree is one of the following four types:

The trees $1\!-\!\overset{1}{\prec}_1$ and $1\!-\!\overset{1}{\prec}_\infty$: The trees $1\!-\!\overset{1}{\prec}_\infty$ can all be exchanged for more of the trees $1\!-\!\overset{1}{\prec}_1$ by the *boundary-twisting* operation [5, p. 1455]. And since $\mathrm{Arf}(L_1) = 0$, the trees $1\!-\!\overset{1}{\prec}_1$ in $t(\mathcal{W})$ have to appear an even number of times [10, 17]. Such trees represent 2-torsion (by antisymmetry), so all these trees can be eliminated from $t(\mathcal{W})$ at the cost of creating only higher-order trees.

The trees t_{2k-1}: Applying the *boundary twist* operation to a clean framed order k Whitney disk $W_{(\cdots(((2,1),1),1),\ldots,1)}$ created by finger moves creates an intersection corresponding to the tree t_{2k-1} at the cost of also creating a twisted Whitney disk whose associated tree is t_k^∞. So, for any k, all t_{2k-1}-trees can be eliminated from $t(\mathcal{W})$, at the cost of only creating new t_k^∞-trees and higher-order trees.

The trees t_{2k}: Applying the *interior twist* operation [5, p. 1456] to a clean order k framed Whitney disk $W_{(\cdots(((2,1),1),1),\ldots,1)}$ created by finger moves creates an intersection corresponding to the tree t_{2k} at the cost of only creating a 2-twisted Whitney disk whose associated tree is t_k^∞. So, for any k, all t_{2k}-trees can be eliminated from $t(\mathcal{W})$, at the cost of only creating new t_k^∞-trees and higher-order trees.

Trees which are zero modulo IHX: As described in [5], using the (twisted) IHX construction, any tree which represents zero modulo IHX relations can be eliminated from $t(\mathcal{W})$ at the cost of creating only higher-order trees.

Thus we can always remove β-bad trees at the cost of creating only higher-order trees, which is sufficient to establish the desired result inductively.

Secondly, we show that the $\beta^i(L)$ are the signed count of t_i^{∞} trees. Decompose the order $2k$ Cochran tower into concordances and simple (twisted) clasper surgeries from the unlink to L, by Theorem 3.1. As shown by Cochran, concordances leave β^i invariant. If a tree is not β-bad, then a corresponding clasper surgery can only change β^i if it is a t_i^{∞}-surgery, by Corollary 3.10, Proposition 3.17 and Proposition 3.16. And by Corollary 3.15, the effect of a t_i^{∞}-surgery is to change $\beta^i(L)$ exactly by $\omega(\Gamma)$. Finally, clasper surgeries whose trees have order $> 2k$ and twisted clasper surgeries whose trees have order $> k$ do not affect β^i for $i \leq k$, by Proposition 3.18. □

Remark 3.19. The construction of a Cochran tower of arbitrary height in the proof of Theorem 1.1 follows from the fact that the Whitney tower obstruction theory allows for β-bad trees to be exchanged for t_i^{∞} trees and higher-order trees. Although we have given examples of how certain β-bad trees can create indeterminacies in the computation of Cochran invariants as the count of twistings on Whitney disks associated to t_i^{∞} trees (Examples 1.11, 1.12 and 2.1), it is possible that some of the trees we have defined to be β-bad might not create indeterminacies in this computation. Eliminating some trees from the β-bad list would *a priori* enlarge the set of links that bound Cochran towers of infinite order, a particularly nice case where all the Cochran invariants can be computed from a single Cochran tower.

Acknowledgments

The authors thank Max-Planck-Institut für Mathematik for its generous hospitality and support. This paper originated there during discussions in Summer 2015 when the first two authors were visiting the third. The second author is also supported by a Simons Foundation *Collaboration Grant for Mathematicians*. The authors would also like to thank Tim Cochran for life-long inspiration.

References

[1] S. Chmutov, S. Duzhin and J. Mostovoy, *Introduction to Vassiliev Knot Invariants* (Cambridge University Press, Cambridge, 2012), xvi+504 pp.

[2] T. Cochran, Geometric invariants of link cobordism, *Comment. Math. Helv.* **60**(2) (1985) 291–311.

[3] T. Cochran, Derivatives of links, Milnor's concordance invariants and Massey products, *Mem. Amer. Math. Soc.* **84**(427) (1990).

[4] J. Conant, R. Schneiderman and P. Teichner, Higher-order intersections in low-dimensional topology, *Proc. Natl. Acad. Sci. USA* **108**(20) (2011) 8131–8138.
[5] J. Conant, R. Schneiderman and P. Teichner, Whitney tower concordance of classical links, *Geom. Topol.* **16**(3) (2012) 1419–1479.
[6] J. Conant, R. Schneiderman and P. Teichner, Milnor invariants and twisted Whitney towers, *J. Topol.* **7**(1) (2014) 187–224.
[7] J. Conant, R. Schneiderman and P. Teichner, Geometric filtrations of string links and homology cylinders, *Quantum Topol.* **7**(2) (2016) 281–328.
[8] J. Conant, R. Schneiderman and P. Teichner, Clasper concordance, Whitney towers and K-repeating Milnor invanants, in preparation.
[9] J. Conant and P. Teichner, Grope cobordism of classical knots, *Topology* **43**(1) (2004) 119–156.
[10] M. Freedman and R. Kirby, A geometric proof of Rochlin's Theorem, *Proc. Sympos. Pure Math.* **32** (1978) 85–98.
[11] M. Freedman and F. Quinn, *The Topology of* 4*-Manifolds*, Princeton Mathematics Series, Vol. 39 (Princeton University Press, Princeton, NJ, 1990).
[12] M. Freedman and P. Teichner, 4-manifold topology. II. Dwyer's filtration and surgery kernels, *Invent. Math.* **122**(3) (1995) 531–557.
[13] M. N. Goussarov, Interdependent modifications of links and invariants of finite degree, *Topology* **37**(3) (1998) 595–602.
[14] K. Habiro, Claspers and finite type invariants of links, *Geom. Topol.* **4** (2000) 1–83.
[15] P. Kirk, Link maps in the four sphere, in *Differential Topology* Lecture Notes in Mathematics, Vol. 1350 (Springer, Berlin 1988), pp. 31–43.
[16] S. Kojima and M. Yamasaki, Some new invariants of links, *Invent. Math.* **54**(3) (1979) 213–228.
[17] Y. Matsumoto, Secondary intersectional properties of 4-manifolds and Whitney's trick, *Proc. Sympos. Pure Math.* **32** (1978) 99–107.
[18] J. Milnor, Link groups, *Ann. of Math.* (2) **59** (1954) 177–195.
[19] N. Sato, Cobordisms of semi-boundary links, *Topology Appl.* **18** (1984) 225–234.