ISSN 1364-0380 (on line) 1465-3060 (printed)

Geometry & Topology
Volume 5 (2001) 551–578
Published: 17 June 2001

Homology surgery and invariants of 3–manifolds

Stavros Garoufalidis
Jerome Levine

School of Mathematics
Georgia Institute of Technology
Atlanta, GA 30332-0160, USA
and
Department of Mathematics
Brandeis University
Waltham, MA 02254-9110, USA

Email: `stavros@math.gatech.edu` and `levine@brandeis.edu`

Abstract

We introduce a homology surgery problem in dimension 3 which has the property that the vanishing of its algebraic obstruction leads to a canonical class of π–algebraically-split links in 3–manifolds with fundamental group π. Using this class of links, we define a theory of finite type invariants of 3–manifolds in such a way that invariants of degree 0 are precisely those of conventional algebraic topology and surgery theory. When finite type invariants are reformulated in terms of *clovers*, we deduce upper bounds for the number of invariants in terms of π–decorated trivalent graphs. We also consider an associated notion of surgery equivalence of π–algebraically split links and prove a classification theorem using a generalization of Milnor's $\bar{\mu}$–invariants to this class of links.

AMS Classification numbers Primary: 57N10

Secondary: 57M25

Keywords: Homology surgery, finte type invariants, 3–manifolds, clovers

Proposed: Robion Kirby
Seconded: Joan Birman, Cameron Gordon

Received: 31 May 2000
Revised: 2 May 2001

1 Introduction

In this paper we take a new approach to the role of finite-type invariants in 3–manifold topology. Our approach is to subdivide the collection of 3–manifold invariants into three increasingly delicate classes.

(1) The invariants of *classical* algebraic topology, which we take to mean invariants of *homology type* in the strongest sense.

(2) Surgery-theoretic invariants.

(3) Invariants which we consider to be of *finite-type*.

In this spirit we start with a base manifold N and then consider all manifolds which are homology equivalent to it, over $\pi_1(N)$. This puts us in a (homology) surgery-theoretic framework with a resulting Witt-type invariant. The vanishing of this invariant then restricts us to a class of manifolds which can be constructed from N by surgery on framed links in N which are algebraically split in a suitable sense. We can, within this class, define a notion of finite-type invariant analogous to earlier notions which were considered, most effectively, for the class of homology spheres (corresponding to $N = S^3$).

A natural question to ask is whether this class of manifolds, and the notion of finite type, depends only on the homology equivalence class of N. We make some progress toward an affirmative answer.

We then give a reformulation of this finite-type theory in terms of what has been recently called Y*–graphs* (see [6]) or *claspers* (see [7]) or *clovers* (see [3]). This will enable us to deduce upper bounds for the number of invariants of a given degree from the number of π–decorated trivalent graphs of that degree.

Finally we consider a notion of surgery equivalence for *algebraically split* links in a general 3–manifold, which is closedly related to our finite-type theory, generalizing the relation between classical surgery equivalence in S^3 and finite-type theory in homology spheres, as explained in [4]. We then define $\mathbb{Z}[\pi \times \pi]$–valued triple Milnor invariants and show that they classify surgery equivalence, generalizing [9]. We also give a direct proof that concordant links are surgery equivalent.

2 A surgery problem

Throughout this paper, all manifolds will be smooth and oriented and all maps will be orientation preserving. Let N be a closed 3–manifold. Consider *degree*

1 maps $f\colon M \to N$, where M is also a closed oriented 3–manifold. Then the induced $f_*\colon \pi_1(M) \to \pi_1(N) = \pi$ is onto and we can consider the induced homomorphism $f_*\colon H_*(\widetilde{M}) \to H_*(\widetilde{N})$, where $\widetilde{N}, \widetilde{M}$ indicates the π–coverings. We will say that f is a $\mathbb{Z}\pi$*–homology equivalence* if f_* (on $H_*(\widetilde{M})$) is an isomorphism and f is degree 1. Since $H_1(\widetilde{N}) = 0$ and $\pi_1(\widetilde{M}) \cong \operatorname{Ker} f_*$ (on $\pi_1(M)$), it follows by Poincaré duality (see [15, Lemma 2.2]) that this is equivalent to the condition that $\operatorname{Ker} f_*$ is a *perfect* subgroup. Given another $\mathbb{Z}\pi$–homology equivalence $f'\colon M' \to N$, we say they are *diffeomorphically equivalent* iff there exists a diffeomorphism $g\colon M \to M'$ such that f is homotopic to $f' \circ g$. Let $\mathcal{H}(N)$ denote the *structure set* of *diffeomorphism equivalence classes* and let $\mathcal{H}_0(N)$ (resp. $\mathcal{H}_0^s(N)$) denote the set of (simple) $\mathbb{Z}\pi$–homology bordism classes of (simple) $\mathbb{Z}\pi$–homology equivalences $f\colon M \to N$.

Our goal in this section is to define a *(homology) surgery obstruction map* Φ and its relatives Φ_0 and Φ_0^s which fit in the following commutative diagram:

Theorem 1

$$\begin{tikzcd}
\mathcal{H}_0^s(N) \arrow[r] \arrow[d, "\Phi_0^s"'] & \mathcal{H}_0(N) \arrow[d, two heads, "\Phi_0"'] & \mathcal{H}(N) \arrow[l, two heads] \arrow[d, two heads, "\Phi"'] \\
\widetilde{W}_s(\pi) \arrow[r] & \widetilde{W}(\pi) & B(\pi) \arrow[l, two heads]
\end{tikzcd}$$

2.1 Algebraic preliminaries

In this section we define the semigroups of equivalence classes of matrices $\widetilde{W}(\pi), \widetilde{W}_s(\pi)$ and $B(\pi)$ over $\mathbb{Z}\pi$ that appear in Theorem 1. These are mild variations of Witt-type constructions, motivated entirely by the geometric results of Section 2.2.

The group-ring $\mathbb{Z}\pi$ has an involution defined by $\overline{ng} = ng^{-1}$ for $n \in \mathbb{Z}$ and $g \in \pi$. Let A be a *Hermitian matrix* over $\mathbb{Z}\pi$, ie, one that satisfies $\bar{A}^t = A$, where ^{t} denotes the transpose. Two Hermitian matrices A, B are *congruent* if there exists a non-singular matrix P such that $B = PA\bar{P}^t$. We say that a Hermitian matrix A is *almost even* if for every $g \in \pi$ with $g^2 = 1$ but $g \neq 1$, the coefficient of g in any diagonal entry of A is even. Note that if A is nonsingular and almost even, so is A^{-1}. Also, any matrix congruent to A is almost even. Given two Hermitian matrices A, B, we will say that they are *stably congruent* if there exist *unidiagonal* matrices S_i such that the block sums $A \oplus S_1$ and $B \oplus S_2$ are congruent. A unidiagonal matrix is a diagonal matrix all of whose diagonal entries are ± 1. Let $\widetilde{W}(\pi)$ denote the set of *stable congruence classes of non-singular almost even matrices.*

Proposition 2.1 *$\widetilde{W}(\pi)$ is an abelian group under block sum.*

Proof We need to show that $A \oplus (-A)$ is stably congruent to a unidiagonal matrix. In fact $A \oplus (-A)$ is congruent to $\left(\begin{smallmatrix} 0 & A \\ A & A \end{smallmatrix}\right)$, which is congruent to $\left(\begin{smallmatrix} 0 & I \\ I & A^{-1} \end{smallmatrix}\right)$ and which, in turn, since A^{-1} is almost even, is congruent to $\left(\begin{smallmatrix} 0 & I \\ I & D \end{smallmatrix}\right)$, where D is a diagonal matrix all of whose entries are 0 or 1. But this is congruent to some unidiagonal matrix. □

Remark 2.2 The proof shows that any *metabolic* non-singular almost even matrix is trivial in $\widetilde{W}(\pi)$. A metabolic matrix is one which is congruent to a matrix of the form $\left(\begin{smallmatrix} 0 & I \\ I & X \end{smallmatrix}\right)$, for some X.

Remark 2.3 As a variation on this we can define A, B to be *simple stably congruent* if $A \oplus S_1 = P(B \oplus S_2)\bar{P}^t$, for some non-singular *elementary* matrix P. An elementary matrix is a product of matrices, each of which differs from the identity matrix in one of the two following ways: (i) there is a single non-zero off diagonal entry, or (ii) one of the diagonal entries is replaced by $\pm g$, for some $g \in \pi$. If we then define $\widetilde{W}_s(\pi)$ to be the set of *simple stable congruence classes of elementary non-singular almost even matrices*, the same proof shows that $\widetilde{W}_s(\pi)$ is a group. There is an obvious homomorphism $\widetilde{W}_s(\pi) \to \widetilde{W}(\pi)$.

Let $B(\pi)$ denote the set of *simple stable congruence classes of almost even non-singular Hermitian matrices.* This is a semigroup under block sum but is not a group since the proof of Proposition 2.1, showing $-A$ is an inverse for A under simple stable congruence, only works if A is elementary—see Remark 2.3. There is an obvious inclusion $\widetilde{W}_s(\pi) \subset B(\pi)$ and epimorphism $B(\pi) \to \widetilde{W}(\pi)$ whose composition $\widetilde{W}_s(\pi) \to B(\pi) \to \widetilde{W}(\pi)$ agrees with the map of Remark 2.3.

2.2 Surgery and a link description of $\mathcal{H}(N)$

It is well-known that the set of closed 3–manifolds can be identified with the set of framed links in S^3 modulo an explicit (Kirby) equivalence relation discussed below. The first goal of this section is to give a similar link description of the set $\mathcal{H}(N)$. The surgery obstruction maps Φ, Φ_0 and Φ_0^s will then be obtained by considering linking matrices of appropriate classes of links.

Lemma 2.4 *If $f\colon M \to N$ is a degree 1 map, then one can adjoin handles of index 2 to N to obtain a compact 4–manifold V with $\partial V - N = M$, and extend f to a map $F\colon V \to N$ so that $F|N =$ identity.*

Proof f represents an element in $\Omega_3(N) \cong H_3(N)$ and so its bordism class is determined by the degree of f. It follows that f is bordant to the identity map of N. This gives us the manifold V and map F without the desired handlebody structure. We must eliminate the handles of index not equal to 2.

First of all we can eliminate handles of index 0 and 4 in the usual way. Now a handle of index 1 represents a boundary connect sum with $S^1 \times D^3$. We can replace this with $D^2 \times S^2$, thereby changing V. The only problem is how to replace the map F on this altered piece. Since F is a retract of N, $F|S^1 \times D^3$ represents an element of $F_*(\pi_1(N))$. By sliding one foot of this 1–handle around a representative of this element in $N \subset V$ we can arrange that $F|S^1 \times D^3$ is null-homotopic. Thus, after replacing this 1–handle with a 2–handle, we can also replace F. To get rid of the 3–handles we regard them as 1–handles on M and apply the same argument since $F_*|\pi_1(M)$ is onto. □

Consider a manifold V as in Lemma 2.4 and a choice of 2–handles. These are attached along a framed link $L \subset N$. Since $F|N = \mathrm{id}$, the components of L are null-homotopic. Conversely, given a framed null-homotopic link L in N we can construct V and then extend the identity map of N over V to obtain a degree 1 map $f\colon N_L \to N$, where N_L denotes the result of surgery on L. The only indeterminacy in this construction is the choice of the extension F. There is no indeterminacy if we assume that N is *prime*, ie $\pi_2(N) = 0$. In case N is not prime we proceed as follows. Consider the set $\mathrm{Deg}(N)$ of all *diffeomorphism classes of degree* 1 maps $f\colon M \to N$. We introduce an *equivalence relation* $\sim$ in $\mathrm{Deg}(N)$ generated by the following modification of a map $f\colon M \to N$. Let $K \subset M$ be a framed knot representing an element in $\mathrm{Ker}\, f_*\colon \pi_1(M) \to \pi_1(N) = \pi$ and let $\phi\colon (S^2, *) \to (N, x_0)$ be any map. The framing of K identifies a neighborhood U of K with $S^1 \times D^2$ and we can assume that $f(U) = x_0$. Now define $f'\colon M \to N$ by $f'|\overline{N-U} = f|\overline{N-U}$ and $f'|U$ is given by the composition

$$S^1 \times D^2 \xrightarrow{p_2} D^2 \xrightarrow{\rho} S^2 \xrightarrow{\phi} N$$

where p_2 is projection on the second factor and ρ is the identification map $(D^2, S^1) \to (S^2, *)$.

The above discussion gives a well-defined onto map $\mathcal{L}(N) \to \widetilde{\mathrm{Deg}}(N)$ of the set $\mathcal{L}(N)$ of *framed, nullhomotopic links*[1] in N to the set $\widetilde{\mathrm{Deg}}(N)$ of $\sim$–equivalence classes of $\mathrm{Deg}(N)$. This map is not one-to-one, since surgery on nonisotopic links may correspond to the same element of $\widetilde{\mathrm{Deg}}(N)$. Recall that two framed links $L, L' \subset N$ are *Kirby equivalent* (which we denote by $L \sim_\kappa L'$) iff N_L is diffeo to N'_L. Fenn–Rourke [1, Theorem 8] showed that Kirby equivalence is generated by three moves (and their inverses) on framed links:

(i) Adding a trivial knot with a ± 1–framing in a ball disjoint from the rest of the link.

(ii) Replacing a component with a connected sum of that component with a push-out of another component, suitably framed.

(iii) Adding a knot with arbitrary framing together with a meridian of it with 0–framing.

In a less advertised part of their paper, Fenn–Rourke considered the *equivalence relation* $\sim_{\kappa'}$ on framed links generated by moves (i) and (ii) alone, see [1, Theorem 6], which is related to surgery on maps rather than surgery on spaces. It is easy to see that $\sim_{\kappa'}$ preserves the class of nullhomotopic links in N and respects the map $\mathcal{L}(N) \to \widetilde{\mathrm{Deg}}(N)$. A converse is given by consequence of a not-so-well known theorem of Fenn–Rourke [1].

Proposition 2.5 *The map $\mathcal{L}(N)/(\sim_{\kappa'}) \to \widetilde{\mathrm{Deg}}(N)$ is one-to-one and onto.*

Proof This follows from [1, Theorem 6]. The isomorphism i in diagram (Δ) of [1] is determined, since the maps $\pi_1(M) \to \pi_1(W(\mathbf{L}_i))$ are isomorphisms. The commutativity of (Δ) corresponds to the diffeomorphism equivalence of the pairs (N, f). $\square$

Given a framed link $L \in \mathcal{L}(N)$, we can define a *linking matrix.* For each component L_i of L choose a lift $\widetilde{L}_i$ in $\widetilde{N}$, the universal cover of N. Then we have linking numbers $\mathrm{lk}(g\widetilde{L}_i, h\widetilde{L}_j)$ for any $g, h \in \pi$—if $i = j$ and $g = h$, we need to push off along some vector field in the given framing. Now the full linking element $\lambda(L_i, L_j) \in \mathbb{Z}\pi$ is defined as $\sum_{g\in\pi} \mathrm{lk}(g\widetilde{L}_i, \widetilde{L}_j)g^{-1}$. We will associate to f the matrix $A = (\lambda(L_i, L_j))$. There is a mild and manageable indeterminacy in the choice of lifts of L. In particular, any change of lifts will change A by a simple congruence.

[1] Not to be confused with the notion of nullhomotopic links in the sense of Milnor.

Proposition 2.6 *The linking matrix A is almost even. If A is any almost even Hermitian matrix over $\mathbb{Z}\pi$ then there is a framed link in N with null-homotopic components whose linking matrix is A.*

Proof Let S be a surface in $\widetilde{N}$ bounded by $\widetilde{L}_i$. If $g \in \pi$ with $g^2 = 1$ but $g \neq 1$, then the coefficient of g in $\lambda(L_i, L_i)$ is the intersection number $S \cdot g\widetilde{L}_i$. Consider the intersection $X = S \cap gS$. This is a collection of loops and arcs in S. $\lambda(L_i, L_i)$ can be computed by counting up (with sign) the number of points of $X \cap \partial S$. Now X is invariant under the action of g and any arc of X with exactly one point on ∂S is sent by g to another such arc in X. Since the action of g is free this means there are an even number of such arcs.

To prove the realizability start out with any null-homotopic link in N. Then the desired matrix A can be obtained from the linking matrix of this link by a sequence of two operations.

(1) For any $g \in \pi$ and i, j add $\pm g$ to a_{ij} and $\pm g^{-1}$ to a_{ji}.

(2) For any i add ± 1 to a_{ii}.

These operations can be realized by changing the link as follows.

(1) Replace L_i by a connected sum of L_i with a small loop linking L_j. The arc used to take the connected sum is determined by g.

(2) Change the framing of L_i by introducing a single twist. □

Remark 2.7 Another interpretation of the linking matrix A is as a representative of the intersection pairing on $\operatorname{Ker}\{F_*\colon H_2(\widetilde{V}) \to H_2(\widetilde{N})\}$. Note that $\operatorname{Ker} F_* \cong H_2(\widetilde{V}, \widetilde{N})$ which is a free $\mathbb{Z}\pi$–module with basis determined by the 2–handles. A is non-singular if and only if f is a $\mathbb{Z}\pi$–homology equivalence and is, in addition, elementary if f is a simple $\mathbb{Z}\pi$–homology equivalence, see Wall [15]. If we now narrow the definition of the relation $\sim$ on $\mathcal{H}(N)$ by restricting our modifications to knots K satisfying $\lambda(K, K) = \pm 1$, in order to stay within the class of $\mathbb{Z}\pi$–homology equivalences, then Proposition 2.5 implies

Proposition 2.8 *There is a one-to-one correspondence*

$$\mathcal{L}^{\mathrm{ns}}(N)/(\sim_{\kappa'}) \to \widetilde{\mathcal{H}}(N)$$

between the set of $\sim$–equivalence classes of $\mathcal{H}(N)$ and the set of $\sim_{\kappa'}$–equivalence classes of $\mathcal{L}^{\mathrm{ns}}(N)$ of framed nullhomotopic links with a nonsingular linking matrix.

Thus, we obtain a well-defined map $\Phi\colon\ \mathcal{H}(N) \twoheadrightarrow B(\pi)$ (which is onto by Proposition 2.6) given by a composition

$$\mathcal{H}(N) \to \widetilde{\mathcal{H}}(N) \cong \mathcal{L}^{\mathrm{ns}}(N)/(\sim_{\kappa'}) \to B(\pi)$$

which assigns to an element $f\colon\ M \to N$ of $\mathcal{H}(N)$ represented by surgery on a framed nullhomotopic link, the linking matrix of that link.

Proposition 2.9 *If $f\colon\ M \to N$ is a $\mathbb{Z}\pi$–homology equivalence then the stable congruence class of A depends only on the $\mathbb{Z}\pi$–homology bordism class of f.*

Proof Suppose that $f'\colon\ M' \to N$ is bordant to the identity on N by $F\colon\ V' \to N$, where V' consists of 2–handles adjoined to N. Let A' be an associated matrix. Suppose f' is bordant to f by a $\mathbb{Z}\pi$–homology bordism $G\colon\ W \to N$. By pasting V, V', W together we create a bordism $\hat{G}\colon\ X \to N$ from the identity map on N to itself. The intersection pairing on $\operatorname{Ker}\hat{G}_*\colon\ H_2(\widetilde{X}) \to H_2(\widetilde{N})$ is represented by $A \oplus (-A')$. Now suppose $\hat{G}$ is bordant, rel boundary, to the projection $I \times N \to N$. Then the standard argument shows that the intersection pairing on $\operatorname{Ker}\hat{G}_*$ is metabolic. Thus by Remark 2.2 the proposition is proved. The obstruction to this bordism is an element of the bordism group $\Omega_4(N) \cong H_4(N) \oplus \Omega_4$. Now $H_4(N) = 0$ and Ω_4 is generated by $\mathbb{C}P^2$ and so this bordism will exist after we connect sum, say, V with a number of copies of $\pm\mathbb{C}P^2$. But this can be achieved by adding to the framed link defining the handlebody decomposition of V a number of trivial components with ± 1–framing. The effect of this is to block sum A with a unidiagonal matrix. $\square$

Thus, we have a well-defined map $\Phi_0\colon\ \mathcal{H}_0(N) \to \widetilde{W}(\pi)$ which is onto by Proposition 2.6. We could also construct an analogous map $\Phi_0^s\colon\ \mathcal{H}_0^s(N) \to \widetilde{W}_s(\pi)$. The commutativity of the diagram of Theorem 1 is obvious.

Example 2.10 Consider $N = S^2 \times S^1, \pi = \mathbb{Z}$. For example if M is obtained by 0–surgery on a knot K in a homology 3–sphere, then there is an obvious degree 1 map $f\colon\ M \to N$, which is a $\mathbb{Z}\pi$–homology equivalence if and only if K has Alexander polynomial 1.

Now suppose $f, g\colon\ M \to N$ are $\mathbb{Z}\pi$–homology equivalences. If $f_* = g_*\colon\ H_1(M) \to H_1(N)$ then it follows from the Hopf classification theorem that $f|M -$ point $\simeq g|N -$ point (homotopic). Thus $f \simeq g\#h$ for some $h\colon\ S^3 \to S^2 \subset S^2 \times S^1$. It follows, using the geometric definition of the Hopf invariant, that $f \sim g$. If $f_* \neq g_*$, then $f_* = g'_*$, where $g' = r \circ g$ and r is the self-diffeomorphism of $S^2 \times S^1$ obtained by reflecting both factors.

This discussion shows that, for any M which is $\mathbb{Z}\pi$–homology equivalent to $N = S^2 \times S^1$, there is either one or two equivalence classes of $\mathbb{Z}\pi$–homology equivalences $M \to N$ depending on whether or not there is an orientation-preserving self-diffeomorphism of M which induces -1 on $H_1(M)$.

2.3 Comparison to surgery theory

We explain here whyc Φ can be thought of, in a rough sense, as encapsulating the *surgery-theoretic* invariants of $\mathcal{H}(N)$. This is not meant to be a mathematically precise statement but more of a philosophical statement.

The surgery exact sequence of Browder–Novikov–Sullivan–Wall extends to lower dimensions in the topological category according to Freedman–Quinn [2]. If N is a closed oriented 3–manifold and $\pi = \pi_1(N)$ is *good* in the sense of Freedman–Quinn (which, admittedly, may be rare) then we have an exact sequence:

$$[\Sigma N : G/\mathrm{Top}] \to L_0^h(\pi) \to \mathcal{H}_0^{\mathrm{top}}(N) \to [N : G/\mathrm{Top}] \to L_3^h(\pi)$$

where $\mathcal{H}_0^{\mathrm{top}}(N)$ is the topological version of $\mathcal{H}_0(N)$ and $L_i^h(\pi)$ are the *Wall surgery obstruction groups* [15]. It is known that $[N : G/\mathrm{Top}]$ can be identified with the set of topological bordism classes of degree 1 maps $M \to N$ where all maps are equipped with a morphism of the tangent bundles. We have isomorphisms $[N : G/\mathrm{Top}] \cong H^2(N; \mathbb{Z}/2)$ and $[\Sigma N : G/\mathrm{Top}] \cong H^1(N; \mathbb{Z}/2) \oplus H^3(N)$, where the $H^3(N) \cong \mathbb{Z}$ summand is mapped isomorphically to $L_0^h(1) \subset L_0^h(\pi)$—see the survey article of Kirby–Taylor [8]. Thus we get an exact sequence

$$H^1(N; \mathbb{Z}/2) \to L_0^h(\pi)/L_0^h(1) \xrightarrow{\tau_h} \mathcal{H}_0^{\mathrm{top}}(N) \to H^2(N; \mathbb{Z}/2) \to L_3^h(\pi) \qquad (1)$$

Similarly we have an exact sequence

$$H^1(N; \mathbb{Z}/2) \to L_0^s(\pi)/L_0^s(1) \xrightarrow{\tau_s} \mathcal{H}_0^{s,\mathrm{top}}(N) \to H^2(N; \mathbb{Z}/2) \to L_3^s(\pi) \qquad (2)$$

for the analogous classification of simple homology equivalence.

There are obvious maps $L_0^h(\pi) \to \widetilde{W}(\pi), L_0^s(\pi) \to \widetilde{W}_s(\pi)$. From the definition of stable congruence these maps induce maps $L_0^h(\pi)/L_0^h(1) \to \widetilde{W}(\pi)$ and $L_0^s(\pi)/L_0^s(1) \to \widetilde{W}_s(\pi)$. According to [14, Prop. 8.2] this is an isomorphism modulo 8–torsion. It is not generally an isomorphism–see [14, Theorem 10.4] for an example with $\pi = \mathbb{Z} \times \mathbb{Z}$.

If π were good, and we were able to ignore the difference between smooth and topological equivalence, we could think of the maps Φ_0 and Φ_0^s as approximations to left inverses of the maps τ_h and τ_s from the sequences (1) and (2).

Thus we can roughly think of $\operatorname{Ker}\Phi$ as those manifolds $\mathbb{Z}\pi$–homology equivalent to N which are undetected by surgery theory. In our theory of finite-type invariants, the invariants of degree 0 will be those which can be recovered from surgery theory, in this sense, and conventional algebraic topology. Those of positive degree can detect differences invisible to surgery theory and conventional algebraic topology.

3 Finite type invariants

3.1 $\mathcal{K}(N)$ and finite type invariants

In this section we study the kernel $\mathcal{K}(N)$ of the map Φ, which leads rather naturally to a distinguished class of links in N and to a notion of finite type invariants.

Suppose $f\colon M \to N$ represents an element of $\mathcal{K}(N)$. Then, by Lemma 2.4, there is a framed link $L \subset N$ with null-homotopic components and linking matrix simply stably congruent to a unidiagonal matrix, such that if V is obtained from N by adding handles along L, then $M = \partial V - N$ and the identity map of N extends to a map $F\colon V \to N$ so that $F|M = f$. Since the moves that define simple stable congruence can be realized by either handle slides, choosing a different lift of L or adding a trivial ± 1–framed component to L we can, in fact, assume that the linking matrix of L is unidiagonal. We will call a (framed) link in N whose components are null-homotopic and linking matrix is unidiagonal *π–algebraically split* (π–AS in short). Conversely, given a π–AS link L in N we can construct V and then extend the identity map of N over V to obtain an element of $\mathcal{K}(N)$. Let $\widetilde{\mathcal{K}}(N)$ denote the set $\mathcal{K}(N)/(\sim)$ for the equivalence relation $\sim$ of Section 2.2 and let $\mathcal{L}^{\mathrm{as}}(N)$ denote the set of framed π–AS links in N. It follows from Proposition 2.5 that $\widetilde{\mathcal{K}}(N)$ is in one–one correspondence with the set $\mathcal{L}^{\mathrm{as}}(N)/(\sim_{\kappa'})$. Note that a handle-slide will usually not preserve the property of being π–AS, so our moves will be sequences of Kirby moves.

We now imitate the usual approach that defines a notion of finite type invariants on a set of objects equipped with a move. Let $\mathcal{F}(N)$ denote the free abelian group on the set $\widetilde{\mathcal{K}}(N)$. For a π–AS link $L \subset N$ we define $[N, L] = \sum_{L' \subset L} (-1)^{|L'|} N_{L'} \in \mathcal{F}(N)$, where N_L denotes the result of surgery of N along L and $|L|$ denotes the number of components of L. Note that the accompanying map $f'\colon N_{L'} \to N$ is uniquely determined by L', modulo $\sim$, so we will suppress f' from the notation. There is a decreasing filtration

$\mathcal{F}(N) = \mathcal{F}_0(N) \supset \mathcal{F}_1(N) \supset \mathcal{F}_2(N)\ldots$ on $\mathcal{F}(N)$, where $\mathcal{F}_n(N)$ denote the subgroup of $\mathcal{F}(N)$ generated by all $[N, L]$ with $|L| \geq n$. We call a function $\lambda\colon \widetilde{\mathcal{K}}(N) \to A$ with values in an abelian group A a *finite type invariant of type* n iff $\lambda(\mathcal{F}_{n+1}(N)) = 0$.

For a decreasing filtration $\mathcal{F}$ (such as $\mathcal{F}(N)$ or $\mathcal{F}^{\mathsf{Y}}(N)$ below) we let $\mathcal{G}$ denote the graded quotients defined by $\mathcal{G}_n = \mathcal{F}_n/\mathcal{F}_{n+1}$.

3.2 Functoriality

Since surgery theory is functorial, we might expect this to also be true of its deformation given by finite type invariants.

Suppose $f\colon N' \to N$ is a $\mathbb{Z}\pi$–homology equivalence between two closed oriented 3–manifolds, where $\pi = \pi_1(N)$. Then $f \in \mathcal{H}(N)$. Consider the induced function $f_*\colon \mathcal{H}(N') \to \mathcal{H}(N)$, where $f_*(g) = f \circ g$, and the further induced function $f_*\colon \widetilde{\mathcal{H}}(N') \to \widetilde{\mathcal{H}}(N)$. We will need the following lemma.

Lemma 3.1 *Suppose K, L are disjoint links in N such that each component of K is null-homotopic in N. Then K is isotopic in N to a link K' such that each component of K' is null-homotopic in $N - L$.*

(1) *If the components of L are also null-homotopic in N, then the linking elements $\lambda(K'_i, L_j) = 0$, for every pair of components of K', L.*

(2) *If K is algebraically split in N, then we can choose K' so that it is algebraically split in $N - L$.*

Proof Since each component K_i of K is null-homotopic in N, it is homotopic in $N - L$ to a product of meridians of L. Thus we can connect sum several meridians of L to K_i to get a new knot K'_i which is null-homotopic in $N - L$ and is clearly isotopic to K_i in the complement of the other components of K.

To see that $\lambda(K'_i, L_j) = 0$, when L_j is null-homotopic, we only need note that any lift $\widetilde{K}'_i$ to the universal cover $\widetilde{N}$ of N is null-homotopic in $\widetilde{N} - \widetilde{L}_j$, for any lift of L_j.

If K is algebraically split then the linking elements $\lambda(K'_i, K'_j) \in \mathbb{Z}\pi$, where $\pi = \pi_1(N - L)$, differ from the entries of a unidiagonal matrix by members of the two-sided ideal I of $\mathbb{Z}\pi$ generated by elements of the form $h - 1$, where $h \in G = \mathrm{Ker}\{\pi_1(N - L) \to \pi_1(N)\}$. Figure 1 shows how to modify K' to change $\lambda(K'_i, K'_j)$ by an element:

(1) $g_1(h-1)g_2$ if $i \neq j$,

(2) $g_1(h-1)g_2 + g_2^{-1}(h^{-1}-1)g_1^{-1}$ if $i = j$

for $g_1, g_2 \in \pi, h \in G$, without changing any other linking element $\lambda(K'_r, K'_s)$ except when $r = j, s = i$.

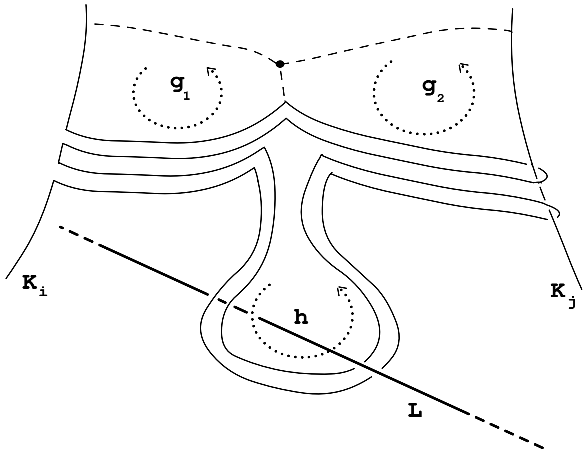

Figure 1: Modification of a link

The dotted curves connecting K'_i, K'_j to the basepoint are those used to specify the lifts—these are needed to define the linking elements. Note that we can represent h by the boundary of a disk in N which is disjoint from K' and the arcs used in the modification, since h is a product of meridians of L. Thus the modified K' is isotopic to K in N.

Since elements of the form (1) generate I, we only need show, by Proposition 2.6, that self-conjugate almost even elements of I are linear combinations of elements of the form (2), which we will call *norm-like*, to conclude that K' can be chosen to be algebraically split.

Choose a subset $S \subset \pi_1(N)$ so that, for every $g \in \pi_1(N)$, exactly one of g, g^{-1} belongs to S. For each $g \in S$ choose $\widetilde{g} \in \pi$ so that $\widetilde{g} \to g$; choose $\widetilde{1} = 1$. Now suppose λ is a self-conjugate even element of I. We can write λ uniquely in the form

$$\lambda = \sum_{g^2 \neq 1} (\lambda_g \widetilde{g} + \widetilde{g}^{-1} \bar{\lambda}_g) + \sum_{g^2 = 1} \lambda_g \widetilde{g} \tag{3}$$

where $\lambda_g \in IG$, the augmentation ideal of $\mathbb{Z}G$. Clearly the terms of the first summation in equation (3) are norm-like, so we consider each term $\lambda_g \widetilde{g}$ of the second summation. Let us write $\lambda_g = \sum_i \epsilon_i g_i$ and so

$$\lambda_g \widetilde{g} = \sum_i \epsilon_i g_i \widetilde{g} \tag{4}$$

where g_i are distinct elements of G, $\epsilon_i \in \mathbb{Z}$ and $\sum_i \epsilon_i = 0$. Since λ is self-conjugate we have $\lambda_g \widetilde{g} = \widetilde{g}^{-1}\bar{\lambda}_g$ and so, for each i there is some j so that $\epsilon_i g_i \widetilde{g} = \epsilon_j \widetilde{g}^{-1} g_j^{-1}$. If $i \neq j$, then replace $\epsilon_j g_j \widetilde{g}$ in equation (4) by $e_i \widetilde{g}^{-1} g_i^{-1}$. If $i = j$ then $g_i \widetilde{g}$ is of order 2 and so ϵ_i is even. In this case rewrite $\epsilon_i g_i \widetilde{g}$ in equation (4), as $2\epsilon_i g_i \widetilde{g}$. Now equation (4) will look like

$$\lambda_g \widetilde{g} = \sum_i \epsilon_i (g_i \widetilde{g} + \widetilde{g}^{-1} g_i^{-1}) \tag{5}$$

where still $\sum_i \epsilon_i = 0$. If we now subtract $0 = \sum_i \epsilon_i (\widetilde{g} + \widetilde{g}^{-1})$ from equation (5) we get

$$\lambda_g \widetilde{g} = \sum_i \epsilon_i ((g_i - 1)\widetilde{g} + \widetilde{g}^{-1}(g_i^{-1} - 1))$$

which is a sum of norm-like terms. □

Proposition 3.2 *We have:*

$$\Phi(f_*(g)) = \Phi(f) + f_*\Phi(g).$$

Proof Suppose $L \subset N$ is an algebraically split link determining $(N' = N_L, f)$ in $\mathcal{H}(N)$ and $K \subset N'$ determines $(M = N'_K, g)$ in $\mathcal{H}(N')$. We can apply Lemma 3.1 to K and the meridians L' of L in N' to allow us to assume that the components of K are null-homotopic in $N' - L' = N - L$. Now $K \cup L$ is algebraically split in N and determines $(M, f \circ g)$ in $\mathcal{H}(N)$. Since, by Lemma 3.1, $\lambda(K_i, L_j) = 0$, the linking matrix of $K \cup L$, which represents $\Phi(f \circ g)$ is the block sum of the linking matrix of L, which represents $\Phi(f)$ and the image under $f_*\colon \pi_1(N') \to \pi$ of the linking matrix of K, which represents $f_*\Phi(g)$. □

It follows from Proposition 3.2 that $f_*(\widetilde{\mathcal{K}}(N')) \subset \widetilde{\mathcal{K}}(N)$.

Proposition 3.3 $f_*(\mathcal{F}_n(N')) \subset \mathcal{F}_n(N)$, *for any* n, *and the induced maps* $f_*\colon \mathcal{G}(N') \to \mathcal{G}(N)$ *are epimorphisms.*

Proof Suppose that $L \subset N$ is an algebraically split link which defines (N', f). Now let $K \subset N' = N_L$ be any algebraically split link—XSwe can assume that K is disjoint from the meridians of L and so lies in $N' - L' = N - L$. In fact, by Lemma 3.1 applied to K and the meridians of L, we can assume that the components of K are null-homotopic in $N-L$ and that $K\cup L$ is algebraically split in N. It is clear that $K\cup L$ defines the element $f_*(M,g)\in\mathcal{K}(N)$.

Suppose that K has n components and so $[N',K]\in\mathcal{F}_n(N')$. Then we have:

$$f_*([N',K])=[N_L,K]=\sum_{L'\subset L}(-1)^{|L'|}[N,K\cup L']$$

Thus we see that $f_*([N',K])\in\mathcal{F}_n(N)$, which shows that $f_*(\mathcal{F}_n(N'))\subset\mathcal{F}_n(N)$, and that $f_*([N',K])\equiv[N,K] \mod \mathcal{F}_{n+1}(N)$, which shows that $f_{*,n}\colon\mathcal{G}_n(N')\to\mathcal{G}_n(N')$ is onto. □

We also have natural filtration-preserving maps $\mathcal{F}(N)\to\mathcal{F}(N\#N')$, for any N,N', defined by

$$f\colon M\to N \quad\rightsquigarrow\quad f\#\mathrm{id}\colon M\#N'\to N\#N'.$$

In particular $\mathcal{F}(N)$ is an $\mathcal{F}(S^3)$ module via a map $\mathcal{F}(S^3)\to\mathcal{F}(N)$, for any N.

4 The $\mathcal{F}(N)$ and $\mathcal{F}^{\mathsf{Y}}(N)$ filtrations

In this section we reformulate our theory of finite type invariants in terms of clovers which in particular allows us to deduce upper bounds for the number of invariants in terms of π–decorated trivalent graphs. Recall the notion of a Y*–graph* in N from [6, 3], the terminology of which we follow here. A Y*–link* is a disjoint union of Y–graphs, and a *clover* is a mild generalization of a Y–link. Given a clover G in N, let N_G denote the result of surgery on G. Throughout this paper, by a Y*–link* or *clover* in a manifold M we will mean one with *nullhomotopic leaves.* This condition, dictated by our surgery problem of Section 2, matches perfectly the generalization of the results of [3] from the case of $N=S^3$ to the case of arbitrary N.

Recall that surgery on a Y–graph is equivalent to surgery on a six component framed link which consists of the three edges and the three leaves of the Y–graph. Since surgery on a Y–graph G (or more generally, a clover) is an example of surgery on a nullhomotopic link with non-singular linking matrix, it follows that $N_G\in\widetilde{\mathcal{H}}(N)$. An alternative geometric proof of this may be obtained from the fact that G lifts to π copies of Y–graphs $\widetilde{G}$ in $\widetilde{N}$ and that the π covering $\widetilde{N_G}$ can be identified with $(\widetilde{N})_{\widetilde{G}}$—since surgery on Y–links in a 3–manifold does not change its homology, it follows that $(\widetilde{N})_{\widetilde{G}}$ is $\mathbb{Z}\pi$–homology cobordant to $\widetilde{N}$. Since the linking matrix of a Y–graph is a metabolic matrix, it follows by Remark 2.2 that $N_G\in\widetilde{\mathcal{K}}(N)$.

Let $\widetilde{\mathcal{K}^{\mathsf{Y}}}(N)$ denote the subset of $\widetilde{\mathcal{K}}(N)$ that consists of all maps N_G for clovers G in N, and let $\mathcal{F}^{\mathsf{Y}}(N)$ denote the free abelian group on $\widetilde{\mathcal{K}^{\mathsf{Y}}}(N)$. We define a decreasing filtration $\mathcal{F}^{\mathsf{Y}}(N) = \mathcal{F}^{\mathsf{Y}}_0(N) \supset \mathcal{F}^{\mathsf{Y}}_1(N) \supset \mathcal{F}^{\mathsf{Y}}_2(N) \dots$ on the abelian group $\mathcal{F}^{\mathsf{Y}}(N)$ where $\mathcal{F}^{\mathsf{Y}}_n(N)$ denotes the span of $[M, G]$ for clovers G in M of degree (ie, number of trivalent vertices) at least n with nullhomotopic leaves.

We will show later that $\widetilde{\mathcal{K}}(N) = \widetilde{\mathcal{K}^{\mathsf{Y}}}(N)$ and that for all integers n, we have $\mathcal{F}^{\mathsf{Y}}_{2n}(N) = \mathcal{F}_{3n}(N)$ after tensoring with $\mathbb{Z}[1/2]$.

4.1 The $\mathcal{A}$–groups and the graded quotients $\mathcal{G}^{\mathsf{Y}}(N)$

The discussion of [3, Section 4.3] implies that $\mathcal{G}^{\mathsf{Y}}_n(N) \otimes \mathbb{Z}[1/2]$ is generated by $[N, G]$ for clovers of degree n without leaves, ie, embedded trivalent graphs of degree n. Unlike the case of $N = S^3$, however, $[N, G]$ depends on the embedding. Consider two embedded trivalent graphs G, G' in N such that $G \smallsetminus e = G' \smallsetminus e'$ for edges e, e' of G, G' which are homotopic, rel. boundary. Then, the Sliding Lemma (in the form of [3, Corollary 4.2]) implies that $[N, G] = [N, G'] \in \mathcal{G}^{\mathsf{Y}}(N)$. The following lemma describes the induced equivalence relation on embedded trivalent graphs.

Lemma 4.1 *For an abstract (not necessarily connected) graph G and path connected space X we have a 1–1 onto map*

$$[G, X] \cong \mathrm{Out}(\pi_1(G), \pi_1(X))$$

where $\mathrm{Out}(\pi_1(G), \pi) = \prod_\alpha \mathrm{Out}(\pi_1(G_\alpha), \pi_1(X))$, *the Cartesian product over the connected components of G, and*

$$\mathrm{Out}(G_1, G_2) = \mathrm{Hom}(G_1, G_2)/(\text{inner automorphisms of } G_2).$$

Proof Pick a *maximal forest* T for G, ie, a maximal tree for each connected component of G. If $f \in [G, X]$, then we can assume that $f(T) = x_0$, a base point of N, ie, that it factors though a map $G/T \to X$. This map is determined by the induced one on the level of π_1. A different choice of a maximal forest or a different choice of a base point of N results in maps on π_1 that differ in each connected component of G by independent inner automorphisms of $\pi_1(X)$. $\square$

Let $\mathcal{A}'(\pi)$ denote the abelian group generated by pairs (G, α') for abstract (not necessarily connected) vertex-oriented trivalent graphs G together with an $\alpha' \in \mathrm{Out}(\pi_1(G), \pi)$, modulo the AS and IHX relations. We call α' a π–*decoration* of G. For each pair (G, α'), pick an arbitrary embedding of G in N

so that the induced map on the fundamental group coincides with α'. Equip the embedding with an arbitrary framing, thus resulting in a clover in N. [3, Lemma 4.4, Corollary 4.5 and Theorem 4.11] shows that this define a map $\Psi_n\colon \mathcal{A}'(\pi) \to \mathcal{G}^{\mathsf{Y}}_n(N)$. The above discussion implies that:

Theorem 2 *For every n, the map $\Psi_n\colon \mathcal{A}'_n(\pi)\otimes\mathbb{Z}[1/2] \to \mathcal{G}^{\mathsf{Y}}_n(N)\otimes\mathbb{Z}[1/2]$ is onto and functorial with respect to $\mathbb{Z}\pi$–homology equivalences.*

Remark 4.2 Over $\mathbb{Q}$, and for $\pi_1(N) = 1$, it is known that the map Ψ is an isomorphism, due to the existence of sufficiently many invariants first constructed by Le–Murakami–Ohtsuki [10].

We now give an alternative description (closely related to π–AS links, see Section 5) of the notion of π–decoration of a graph. This description generalizes to a decoration of the edges by elements of an arbitrary ring with involution, see Definition 4.4 below.

Consider pairs (G,α) of abstract, vertex-oriented, *edge-oriented* trivalent graphs G, together with a map $\alpha\colon \mathrm{Edge}(G) \to \mathbb{Z}\pi$ that colors each oriented edge of G by an element of $\mathbb{Z}\pi$. Let $\mathcal{A}(\pi)$ denote the abelian group generated by pairs (G,α) modulo the relations shown in Figure 2.

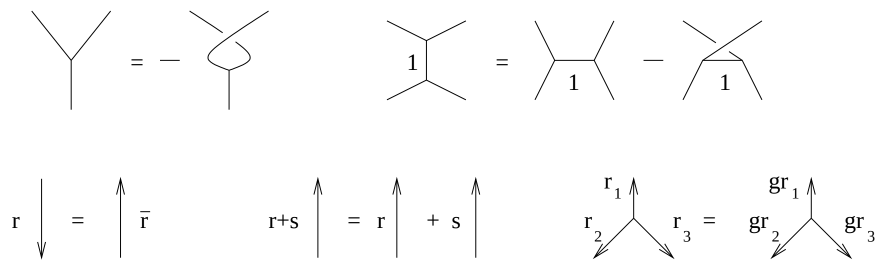

Figure 2: The AS, IHX, R_1, R_2 and R_3 relations. Here $\bar{g} = g^{-1}$ is the involution of $\mathbb{Z}\pi$, $r, s, r_i \in \mathbb{Z}\pi$ and $g \in \pi$.

Lemma 4.3 *There is an isomorphism*

$$\mathcal{A}(\pi) \to \mathcal{A}'(\pi).$$

Proof It suffices to consider a vertex-oriented, edge-oriented *connected* graph G. To a map α, we will associate a map α' and vice versa.

Given a map $\alpha\colon \mathrm{Edge}(G) \to \pi$, (which in view of relation R_2 we may assume that it is a decoration of the edges of G by elements of π) we define a map

$\alpha'\colon \pi_1(G) \to \pi$ as follows. For a closed path of oriented edges $e = (e_1, \ldots, e_n)$, we set $\alpha'(e) = \alpha(e_1)\ldots\alpha(e_n)$. It is easy to see that this defines a group homomorphism $\pi_1(G) \to \pi$, compatible with the relations R_1 and R_3. Conversely, given a map α', choose a maximal tree T and define $\alpha(\mathrm{Edge}(T)) = \{1\}$. Since $\pi_1(G)$ can be identified with the free group on $\mathrm{Edge}(G \smallsetminus T)$, α' will then determine α on these edges. □

The following concept of the $\mathcal{A}$–*groups*, motivated by Theorem 2, has several applications which will be presented in a later publication.

Definition 4.4 Given a ring R with involution and a subgroup U of its group of units, we define $\mathcal{A}(R, U)$ to be the graded abelian group generated by trivalent graphs (with a vertex and an edge orientation) whose edges are decorated by elements of R, modulo the relations of Figure 2, with $r, s, r_i \in R$ and $g \in U$.

4.2 The equivalence of the $\mathcal{F}(N)$ and $\mathcal{F}^{\mathsf{Y}}(N)$ filtrations

In [3, Sections 5.2–5.6] it was shown that finite type invariants of integral homology 3–spheres based on surgery on AS links coincide with those based on surgery on clovers. In this section we will extend this to 3–manifolds, by using the same idea as in [3, Sections 5.2–5.6], together with Lemma 4.5 and Proposition 4.6.

Recall that the proofs of [3, Sections 5.2–5.6] consist of three types of arguments:

- First, *untying* AS–links by clovers and vice versa.
- Second, counting arguments.
- Third, an application of the topological calculus of clovers to the study of $\mathcal{G}^{\mathsf{Y}}(N)$.

The second type of arguments works without change when we replace S^3 by N. So does the third type of argument, since, restricted to clovers with nullnomotopic leaves, it uses the moves (i) and (ii) of Kirby equivalence and the move (iii) for nullhomotopic knots, which is a consequence of (i) and (ii) as shown in [1, Theorem 2]. The first type of argument requires some additional work, which consists of Lemma 4.5 and Proposition 4.6.

Theorem 3 *For all integers n we have that $\mathcal{F}_{3n}(N) \otimes \mathbb{Z}[1/2] = \mathcal{F}^{\mathsf{Y}}_{2n}(N) \otimes \mathbb{Z}[1/2]$.*

Together with Theorem 2, it implies that:

Theorem 4 *For all integers n there is an onto map*

$$\mathcal{A}'_{2n}(\pi) \otimes \mathbb{Z}[1/2] \twoheadrightarrow \mathcal{G}_{3n}(N) \otimes \mathbb{Z}[1/2]$$

functorial with respect to $\mathbb{Z}\pi$–homology equivalences.

4.3 Undoing clovers by AS–links and vice versa

Let G be a Y–link in N. Using the terminology of [3, Section 5.3], we say that a link O in $N \smallsetminus G$ *laces* G, if O is trivial, unimodular and each of the (pairwise disjoint) discs bounding its components intersects G in at most two points, which belong to the leaves of G. G is *trivial*, if it consists of n Y–graphs, standardly embedded in n disjoint balls which lie in an embedded ball in N.

Lemma 4.5 *[3, Lemma 5.3] Let T be a trivial n–component Y–link in N. For any n–component Y–link G in N, there exists a unimodular link O in N which laces T, such that $[N,G]=[N_O,T]$. Under surgery on $T' \subset T$, O gets transformed to a π–AS link in N.*

A Y–link G in $M \smallsetminus L$ *laces a link* L if L is π–AS, G has nullhomotopic leaves and every leaf l of G either bounds a disk which intersects L geometrically once, or the equivariant linking number of l and every component of L vanishes. We call (G,L) a *lacing pair*. A special case of a lacing pair (G,L) for a trivial Y–link G was called a *Borromean surgery* in [11] and a Δ*–move* in [12].

Proposition 4.6 *Let O be a trivial unimodular n–component link in N. For any n–component π–AS–link L in N, there exists a lacing pair (G,O) such that O is trivial unimodular, under surgery on G (N,O) gets transformed into (N,L) and under surgery on $O' \subset O$, G gets transformed to a Y–link in N with nullhomotopic leaves.*

Proof Choose a base point x_0 of N and a *basing* γ of L, ie, a choice of disjoint paths $\{\gamma_i\}$ in $N \smallsetminus L$ from x_0 to points $x_i \in L_i$, one for each component of L. Choose a framing of L and a lift $\widetilde{x}_0 \in \widetilde{N}$ of x_0. Then, there is a unique lift $\widetilde{L} \cup \widetilde{\gamma}$ in $\widetilde{N}$ of $L \cup \gamma$ that contains $\widetilde{x}_0$ and a well-defined linking matrix A of L.

Choose a regular homotopy L^t from $L^0 = L$ to $L^1 = O$, which we can assume is stationary on every x_i. We may assume that L^t is a link except for finitely many times $\{t_s\}$ where L^{t_s} is an immersion with a single transverse double point. The linking matrix $A_{-\epsilon}$ and A_ϵ of $L^{t_s-\epsilon}$ and $L^{t_s+\epsilon}$ are related as follows: if the double point p involves the components L_i and L_j of L^{t_s}, with

$i \leq j$, construct a loop g_p by starting at x_0, going along γ_i to L_i, then over to L_j and back to x_0 via $\overline{\gamma}_j$—see the figure below.

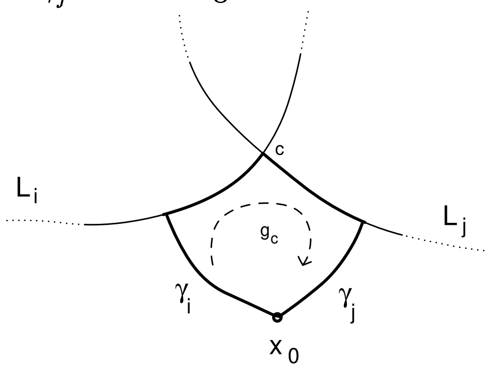

It is easy to see that

$$A_\epsilon = A_{-\epsilon} + \epsilon_p(g_p E_p + g_p^{-1} E_p^t)$$

where $\epsilon_p \in \{-1, +1\}$ is the local orientation sign of p and E_p is the matrix with all zeros except in the (i,j) place where it equals 1. Thus, if A_L is the linking matrix of L and A_O is the linking matrix of an unlink with the same framing and number of components as L, we have $A - A_O = \sum_p \epsilon_p(g_p E_p + g_p^{-1} E_p^t)$, where the sum is over all double points of the homotopy. Since L is a π–AS link, it follows that for every pair of components L_i and L_j of L there is a pairing of the double points of L_i and L_j into classes (p^+, p^-) such that $g_{p^+} = g_{p^-}$ and $\epsilon_{c^+} = -\epsilon_{c^-}$. In other words, one can undo L by a sequence of *double crossing changes* shown below in Figure 3, where the loop $g \simeq g_{p^+} g_{p^-}^{-1} = 1 \in \pi$

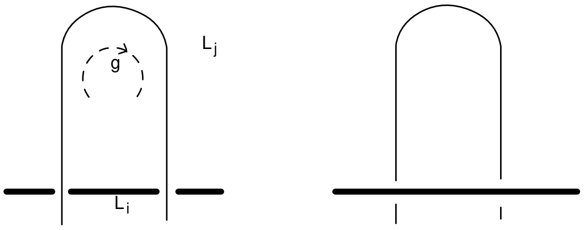

Figure 3: A double crossing change

is nullhomotopic. These double crossing changes can be achieved by surgery on Y–graphs whose leaves are nullhomotopic, see [11, 12] and also Lemma 4.7 below. So far, each of the Y–graphs have two leaves that bound a disk that intersects L at most once and a nullhomotopic leaf. Observe that every nullhomotopic leaf in N bounds a disk with clasp intersections as shown below.

Using repeatedly Move Y_4 of [3, Theorem 3.1] (the so-called, move of *Cutting a Leaf*), as follows

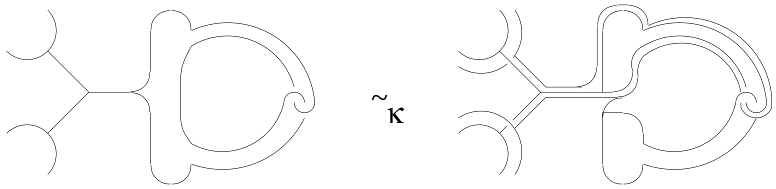

we may assume that every leaf of each Y–graph bounds a disk that either intersects L geometrically once, or none. In all cases, the Y–link G that consists of all these Y–graphs is lacing the unlink O. It is easy to verify that the rest of the statements of the proposition. □

The following lemma shows how to slide a band of an embedded surface through any collection of bands (or leaves of Y–graphs) by Y–surgery.

Lemma 4.7 *The following framed links are Kirby equivalent:*

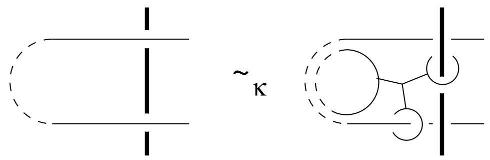

In particular, a double crossing move can be obtained by surgery on a Y*–graph.*

Proof

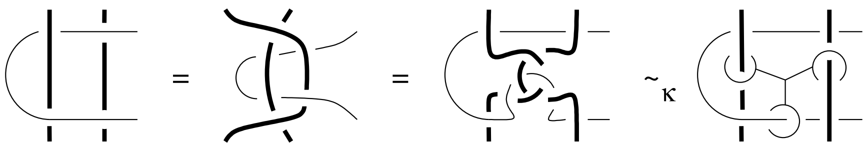

□

Proposition 4.8 $\widetilde{\mathcal{K}}(N) = \widetilde{\mathcal{K}^{\mathsf{Y}}}(N)$.

Proof Since $\widetilde{\mathcal{K}^{\mathsf{Y}}}(N) \subset \widetilde{\mathcal{K}}(N)$, we need only show the opposite inclusion. Proposition 4.6 implies that for every π–AS link L in N, there exists a trivial Y–link T that ties a trivial unimodular link O such that $(N, L) = (N_T, O)$. Let G denote the image of $T \subset N$ under surgery on O. G is a Y–link (with nullhomotopic leaves) and $N_L = N_{T \cup O} = N_G$. The result follows. □

5 Surgery equivalence

In this section we discuss the notion of surgery equivalence of π–AS links, motivated both by surgery theory and by the theory of finite type invariants.

Suppose that L is an unframed π–AS link in N. Now let L, with some unit framing, be expanded to a π–AS link $L \cup L_{\text{triv}} \subset N$ for some trivial, unit-framed link L_{triv}. Let L' denote the image of L under the obvious isomorphism $N \cong N_{L_{\text{triv}}}$. *Surgery equivalence* is the relation on the set of unframed π–AS links in N generated by the move that replaces L by L' for some link L_{triv} as above.

It was shown in [9] that, when $N = S^3$, surgery equivalence classes of unframed AS–links are determined by the Milnor triple $\bar{\mu}$–invariants. We will generalize the construction of these $\bar{\mu}$–invariants in an equivariant manner to define surgery equivalence invariants of π–AS links $L \subset N$. Choose a lift $\widetilde{L}$ of L in the π–cover $\widetilde{N}$. The components $\widetilde{L}_i$ of $\widetilde{L}$ bound oriented surfaces $V_i \subset \widetilde{M}$ and, since L is algebraically split, we can assume that the interior of V_i does not intersect $p^{-1}(L)$, where $p\colon \widetilde{N} \to N$ is the projection. For any $1 \leq i,j,k \leq q$ and $g,h \in \pi$ we define $\bar{\mu}^{g,h}_{ijk}(L)$ to be the triple intersection number of V_i, gV_j, hV_k, when these are three different surfaces. This is independent of the choice of $\{V_i\}$. A change in the choice of liftings $\{\widetilde{L}_i\}$ produces the following change in $\bar{\mu}^{g,h}_{ijk}(L)$: for any $g_1, \dots, g_n \in \pi$ we have

$$\text{new } \{\bar{\mu}^{g,h}_{ijk}(L)\} = \text{old } \{\bar{\mu}_{ijk}^{g_i^{-1}gg_j,\; g_i^{-1}hg_k}(L)\}$$

Note that for the special case when π is abelian, there is no indeterminacy in $\bar{\mu}^{g,h}_{iii}(L)$.

Also note the following:

(1) $\bar{\mu}^{g,h}_{ijk}(L) = \bar{\mu}^{g^{-1}h,g^{-1}}_{jki}(L), \quad \bar{\mu}^{h,g}_{ikj}(L) = -\bar{\mu}^{g,h}_{ijk}(L)$.

(2) $\bar{\mu}^{g,h}_{ijj}(L)$, $(i \neq j)$ is defined only if $g \neq h$.

(3) $\bar{\mu}^{g,h}_{iii}(L)$ is defined only if $1, g, h$ are distinct.

Now set

$$\bar{\mu}_{ijk}(L) = \sum_{g,h\in\pi} \bar{\mu}^{g,h}_{ijk}(L)(g,h) \in \mathbb{Z}[\pi \times \pi]$$

This is a finite sum. In cases (2) and (3), the relations in (1) impose conditions on $\bar{\mu}_{ijj}(L)$. In case (2) it is skew-symmetric in the two factors, but for case (3) the constraint is more complicated.

Example 5.1 If $\pi = \mathbb{Z}$, then $\bar{\mu}_{iii}(L)$ is determined by the $\{\bar{\mu}^{s,t}_{iii}(L) \mid s > t > 0\}$. As an example of a knot K with $\bar{\mu}^{s,t}_{111} = 1$, choose a Borromean link L_1, L_2, L_3 in a ball in N and let K be the connected sum $L_1 \# L_2 \# L_3$, where the tube connecting L_1 to L_2 winds t times around the generator of π and the tube connecting L_2 to L_3 winds $s-t$ times around. Note that $\bar{\mu}^{s',t'}_{111}(K) = 0'$ unless $s = s'$ and $t = t'$. See the figure below for $s = t = 1$.

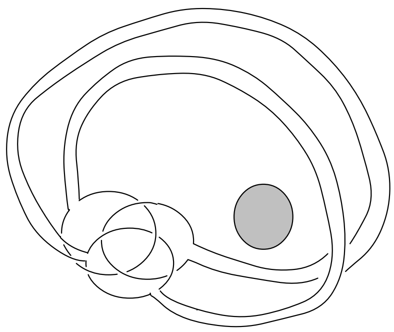

We will show that, just as in the case of unframed links in S^3, the surgery equivalence class is determined by the $\mu^{g,h}_{ijk}$. First we need the following

Lemma 5.2 *If two links L and L' differ as in the first two frames of this picture*

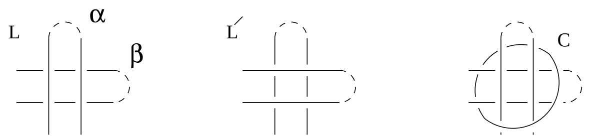

or the first two frames of this picture

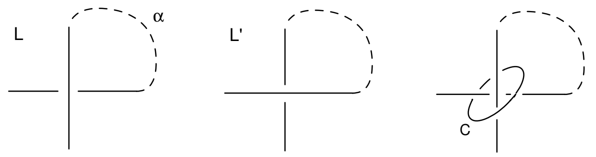

where α, β are nullhomotopic paths, then they are surgery equivalent.

Proof This follows by elementary properties of Kirby's calculus applied to the unit-framed knots C shown in the pictures above. □

Theorem 5 *Suppose L and L' are unframed π–AS links in N. Then L is surgery equivalent to L' if and only if $\mu^{g,h}_{ijk}(L) = \mu^{g,h}_{ijk}(L')$ for all i, j, k, g, h for which they are defined.*

Proof First of all we prove the invariance of the $\mu^{g,h}_{ijk}(L)$ under surgery equivalence. If L_{triv} is chosen so that $L \cup L_{\mathrm{triv}}$ is π–AS, then we can, by tubing, arrange that the surfaces $\{V_i\}$ used to define $\mu^{g,h}_{ijk}(L)$ are disjoint from the lifts of L_{triv} and so pass unchanged into the π–covering of the surgered link. In particular the intersections which define $\mu^{g,h}_{ijk}(L)$ are unchanged.

Now suppose that L and L' are two π–AS links such that $\mu^{g,h}_{ijk}(L) = \mu^{g,h}_{ijk}(L')$. By Proposition 4.6 we know that L can be transformed into L' by surgery on a set of Y–links whose leaves are meridians of L. Since surgery on such a Y–link is the same as a sequence of disjoint Δ*–moves* in the terminology of [12]—see Figure 4—it is easy to see the effect of such a surgery on the $\{\mu^{g,h}_{ijk}\}$. Suppose a

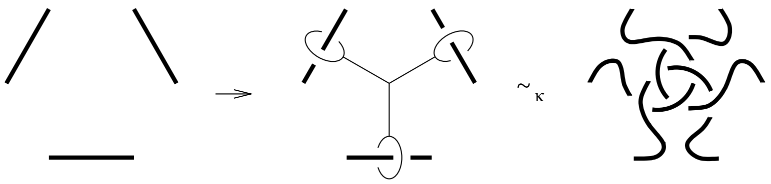

Figure 4: A Δ–move

lift of the Y–link G into $\widetilde{N}$ has its meridians on three components $\widetilde{L}_i, g\widetilde{L}_j, h\widetilde{L}_k$, where $i \le j \le k$. If any two of these components are the same then there is no change in any of the $\mu^{g',h'}_{i'j'k'}$. If the three components are distinct then $\mu^{g,h}_{ijk}$ is changed by ± 1 and every other $\mu^{g',h'}_{i'j'k'}$, where $i' \le j' \le k'$, is unchanged. Thus our assumption about L, L' says that the transformation from L to L' is accomplished by a sequence of surgeries of two types:

(i) surgery on pairs of Y–links $\{G_i, G'_i\}$, where G_i and G'_i can be lifted to Y–links in $\widetilde{N}$ with oppositely oriented trivalent vertices and which have leaves on the same three distinct components, and

(ii) surgeries on individual Y–links G_j with at least two leaves on the same component.

In case (ii) it is easy to see that surgery on G_j does not change the surgery equivalence class since we can undo the Borromean part of the Y–link by crossing changes, using the second part of Lemma 5.2, on the two rings attached to the same component of L.

Thus it remains to show that the effect of surgery on a pair of Y–links G, G' with leaves on the same three distinct components of $\widetilde{L}$ does not change the surgery equivalence class of L. First of all we can consider the case where G is an inverse of G' in the sense of [3, Theorem 3.2]. In this case a surgery on G and G' does not change L at all. For any other G we can assume that there is a homotopy in N from G to an inverse of G' which is stationary on the leaves of G. Such a homotopy is a sequence of isotopies in $N - L$ together with (i) crossings of an edge of G and a component of L and (ii) crossings of an edge of G with a leaf of G. It suffices to show that these two types of crossings do not change the surgery equivalence class of the G surgery on L.

For (i) the effect of this crossing on surgery of L is pictured in Figure 5.

Figure 5: A link component crossing an edge of a Y–graph

The surgery equivalence is given by Lemma 5.2 and the double crossing change in Figure 6.

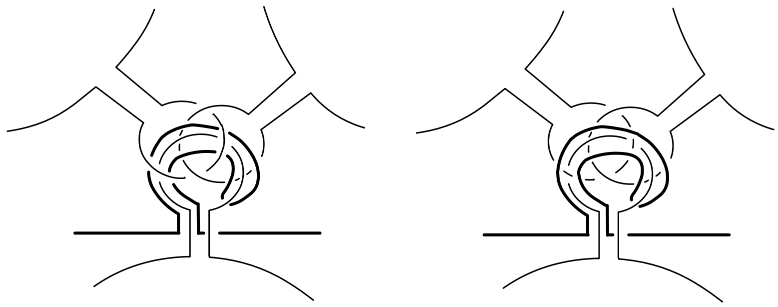

Figure 6: A double crossing change which implements the crossing of Figure 5

For (ii) we invoke the following:

Lemma 5.3 *Suppose G is the union of two* Y*–links G_1, G_2 in the complement of a link L, whose leaves are meridians of L, and G' is obtained from G by a single crossing change of a leaf of G_1 with a leaf of G_2. Then the link produced by surgery on L using G is surgery equivalent to the link produced by surgery on G'.*

Proof By [3, Theorem 2.3] and Y_4 moves of [3], surgery on G' is the same as surgery on G together with surgery on a clover of degree 2 with the shape of 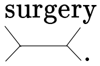. Thus we need to see that surgery on such clovers does not change the surgery equivalence class of L. The effect of surgery on such a clover is shown in Figure 7.

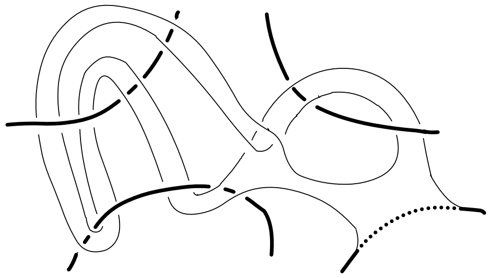

Figure 7: The effect of surgery on a clover of degree 2

A double crossing change which will undo this surgery is illustrated in Figure 8. □

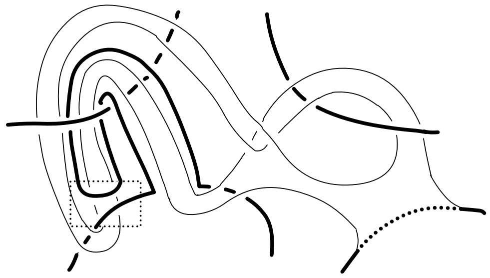

Figure 8: The doublecrossing change inside the box will undo the surgery.

This completes the proof of Theorem 5. □

We also prove that:

Theorem 6 *Concordant links are surgery equivalent.*

Proof Concordance, just as in the classical case, is generated by the following *ribbon move* $L \to L'$. Given a π–AS link $L \subset N$, consider also a finite number of disks $\{D_i\}$ in N, disjoint from each other and L. For each D_i choose a band B_i connecting D_i to a component, which we denote L_i, of L. The band cannot intersect L or any ∂D_j except at its ends. Then L' is defined to be the band-sum of L with $\bigcup \partial D_i$.

Now choose some place where a band B_i penetrates a disk D_j. Choose a path γ from L_j to nearby the penetration so that the closed path consisting of γ followed by the path from D_j along B_j and back along L_j to the starting point of γ is null-homotopic. Now thicken γ to a band (or finger) and apply Lemma 5.2 as follows:

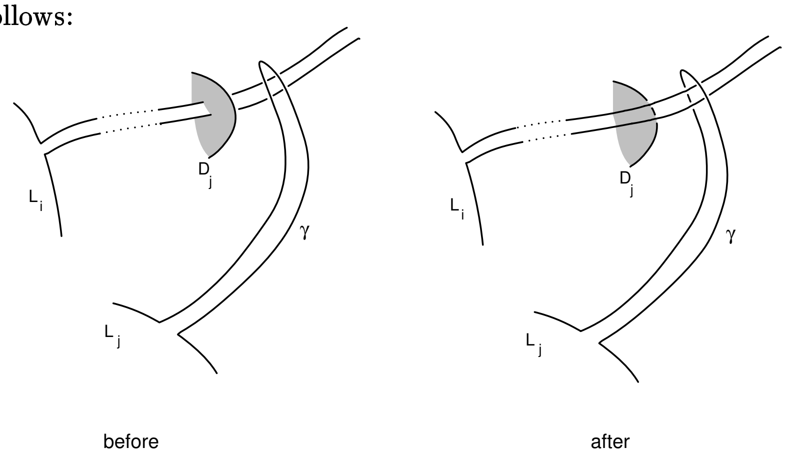

This removes the penetration. Eventually we can remove all the penetrations and the resulting link will be isotopic to L. □

5.1 An alternative study of $\mathcal{G}(N)$.

In this section we mention, in brief, an alternative study of $\mathcal{G}^{\mathsf{Y}}(N)$ using our results on surgery equivalence and the group $\mathcal{A}(\pi)$ from Section 4.1. For $N = S^3$ this coincides with the approach to finite type invariants introduced by Ohtsuki [13], and studied in [5].

Recall the map $\mathcal{L}^{\mathrm{as}}(N) \to \widetilde{\mathcal{K}}(N)$ defined by doing surgery on a unit-framed π–AS link in N. If L and L' are surgery equivalent n component π–AS links in N, then $[N, L] = [N, L'] \bmod \mathcal{F}_{n+1}(N)$. Thus $\mathcal{G}(N)$ is a quotient of the free abelian group on the set of surgery equivalence classes of π–AS links. For each generator (G, α) of $\mathcal{A}(\pi)$ we can define a π–AS link as follows. First construct a link in a 3–ball $B \subset N$ associated to G, using the construction of Ohtsuki [13], by banding together copies of the Borromean rings, one Borromean rings for each vertex and one band for each edge. Now for each edge e of G pull the band corresponding to e around a loop representing $\alpha(e)$, using the orientation of e to direct it. See Figure 9 below. The surgery equivalence class of this link

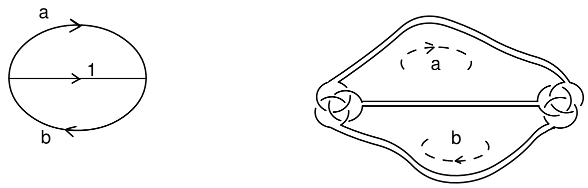

Figure 9: An edge and vertex oriented diagram and the associated π–AS link

$L(G)$ is well-defined since we only have to worry about band-crossings which are covered by Lemma 5.2. Note that if G has degree $2n$, then $L(G)$ has $3n$ components. Using a local 3*–band relation* of [13, Lemma 4.1] for an arbitrary manifold N and the above discussion, we obtain an onto map from the abelian group generated by pairs (G,α) of degree$(G) = 2n$ to $\mathcal{G}_{3n}(N) \otimes \mathbb{Z}[1/2]$. The work of [5], formulated for arbitrary manifolds rather than S^3, implies that the AS and IHX relations of Figure 2 are satisfied, thus obtaining an onto map $\mathcal{A}_{2n}(\pi) \twoheadrightarrow \mathcal{G}_{3n}(N) \otimes \mathbb{Z}[1/2]$ for every integer n. It is an easy exercise in Kirby calculus to show that the above maps fit in a commutative diagram

$$\begin{tikzcd}
\mathcal{A}_{2n}(\pi) \arrow[r, two heads] \arrow[d, leftrightarrow, "\cong"] & \mathcal{G}_{3n}(N) \otimes \mathbb{Z}[1/2] \arrow[d, leftrightarrow, "\cong"] \\
\mathcal{A}'_{2n}(\pi) \arrow[r, two heads] & \mathcal{G}^{Y}_{2n}(N) \otimes \mathbb{Z}[1/2]
\end{tikzcd}$$

Acknowledgement The authors were partially supported by NSF grants DMS-98-00703 and DMS-99-71802 respectively, and by an Israel-US BSF grant.

References

[1] **R Fenn**, **C Rourke**, *On Kirby's calculus of links*, Topology, 18 (1979) 1–15

[2] **M Freedman**, **F Quinn**, *Topology of 4–manifolds*, Princeton University Press, Princeton, NJ (1990)

[3] **S Garoufalidis**, **M Goussarov**, **M Polyak**, *Calculus of clovers and finite type invariants of 3–manifolds*, Geometry and Topology, 5 (2001) 75–108

[4] **S Garoufalidis**, **J Levine**, *Finite-type invariants of 3–manifolds II*, Math. Annalen 306 (1996), 691–718

[5] **S Garoufalidis**, **T Ohtsuki**, *On finite type 3–manifold invariants III: manifold weight systems*, Topology, 37 (1998) 227–244

[6] **M Goussarov**, *Finite type invariants and n–equivalence of 3–manifolds*, CR Acad. Sci. Paris Ser. I Math. 329 (1999) 517–522

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

[8] **R Kirby**, **L R Taylor**, *A survey of 4–manifolds though the eyes of surgery*, from: "Surveys on surgery theory", vol. 1, ed. S Cappell, A Ranicki, J. Rosenberg, Annals of Math. Studies 145, Princeton U. Press (2000)

[9] **J Levine**, *Surgery equivalence of links*, Topology, 26 (1987) 45–61

[10] **T T Q Le**, **J Murakami**, **T Ohtsuki**, *A universal quantum invariant of 3–manifolds*, Topology, 37 (1998) 539–574

[11] **S V Matveev**, *Generalized surgery of three-dimensional manifolds and representations of homology spheres*, Math. Notices Acad. Sci. USSR, 42:2 (1987) 651–656

[12] **H Murakami**, **Y Nakanishi**, *On a certain move generating link homology*, Math. Annalen, 284 (1989) 75–89

[13] **T Ohtsuki**, *Finite type invariants of integral homology 3–spheres*, J. Knot Theory and its Ramifications, 5 (1996) 101–115

[14] **A Ranicki**, *The algebraic theory of surgery I. Foundations*, Proc. London Math. Soc. 40 (1980) 87–192

[15] **C T C Wall**, *Surgery on manifolds*, Academic Press, London (1970)