LECTURES ON EXACT TRIANGLES IN HEEGAARD FLOER THEORY

IAN ZEMKE

Ozsváth and Szabó proved a surgery exact triangle:

$$\cdots \to \widehat{HF}(Y) \to \widehat{HF}(Y_\lambda(K)) \to \widehat{HF}(Y_{\lambda+1}(K)) \to \cdots$$

for Heegaard Floer homology. Our goal is to understand its proof and a few applications. We will also discuss the knot surgery formula, the algebra of L-space knots and links, and finally discuss briefly Heegaard Floer homology for plumbed 3-manifolds.

Ozsváth and Szabó have written additional introductory surveys, which are also an excellent resource and overlap with some of the below topics. See
`https://web.math.princeton.edu/~petero/`
to find these surveys.

1. Lecture 1: Type-D structures and twisted complexes

1.1. **Dehn surgery.** Dehn surgery is one of the most fundamental operations in 3-manifold topology. Given a knot K, we remove a regular neighborhood of K, denoted $\nu(K) \cong S^1 \times D^2$, and then reglue it using a diffeomorphism of $\partial S^1 \times D^2 = \mathbb{T}^2$.

Note that gluing $S^1 \times D^2$ can be accomplished by gluing in two steps. View S^1 as $[0,1]/(0 \sim 1)$. First glue in $[0,1/2] \times D^2$ (a thickened disk). For this, we only need to know where $\lambda := \{0\} \times D^2$ is sent, which we can encode as a simple closed curve on $\mathbb{T}^2 = \partial S^3 \setminus \nu(K)$. We call this a *framing* of the knot.

Note that after gluing in $\{0, 1/2\} \times D^2$, we are left with a manifold with boundary S^2. Gluing in $\{1/2, 1\} \times D^2 \sim B^3$ amounts to gluing B^3 to the resulting boundary. Up to homeomoprhism, the gluing map for $\{1/2, 1\} \times D^2$ (after having already glued $\{0, 1/2\} \times D^2$) does not affect the resulting 3-manifold up to homeomorphism by the Alexander trick (see Exercise 1.1). In particular, the only data needed to determine the surgery is the longitude.

If λ is a choice of framing, we write $Y_\lambda(K)$ for λ framed surgery.

Exercise 1.1. Prove that any homeomorphism $f\colon S^2 \to S^2$ extends to a homeomorphism of B^3 with itself. (In fact, your argument should in fact show that any self homeomorphism of S^n extends to a self-homeomorphism of B^{n+1}). This is typically called the *Alexander trick*.

Remark 1.2. The orientation of K does not affect the Dehn surgery $Y_\lambda(K)$.

In these lectures, we focus on the case that λ is a *Morse framing* (also called a *longitudinal framing*), which is when the map

$$H_1(\mathbb{T}^2) \to H_1(\nu(K))$$

maps λ to a generator of $H_1(\nu(K))$. Typically if K is oriented, we will require the longitude to map to the corresponding generator of $H_1(\nu(K))$.

Note that given an orientation of K, there is a canonical meridian $\mu \in H_1(\partial\nu(K))$, whose orientation is given by the right hand rule. Hence, given an oriented knot K with a Morse framing λ (chosen compatibly with the orientation of K), there is another framing $\lambda+1$ obtained by adding μ to λ.

Exercise 1.3. Show that although the choice of orientation does affect the meridian of K, the choice of orientation does not affect the surgery $Y_{\lambda+1}(K)$.

Remark 1.4. We often say that Y itself is the result of "∞-framed" surgery on K (cutting out a neighborhood of K and regluing using the identity map). The reason is because there is a more general operation called *rational surgery.* Given a knot K in S^3, there is a preferred longitude, called the *Seifert longitude*, which spans the kernel of the map $H_1(\partial\nu(K)) \to H_1(S^3 \setminus \nu(K))$. Given p and q coprime, we define the p/q framed surgery as Dehn surgery where the longitude intersects the Seifer longitude p times and the meridian q times. With this notation, regluing $\nu(K)$ using the identity corresponds to doing $1/0 = \infty$ framed surgery.

Dehn surgery is a fundamental operation in 3-manifold theory. The following theorem tells us why:

Theorem 1.5 (Lickorish-Wallace 60's)**.** *Every closed, oriented 3-manifold can be obtained by Dehn surgery on a link in S^3.*

1.2. **Statement of the exact triangle.** In this lecture, we are interested in proving the surgery exact triangle of Ozsváth and Szabó:

Theorem 1.6. *If K is a knot in Y with longitudinal framing λ, then there is an exact triangle*

$$\cdots \to \widehat{HF}(Y) \to \widehat{HF}(Y_\lambda(K)) \to \widehat{HF}(Y_{\lambda+1}(K)) \to \cdots .$$

The exact triangle also holds for HF^- if we complete by tensoring with the power series ring $\mathbb{F}[[U]]$.

Ozsváth and Szabó prove the exact triangle by exhibiting a homotopy equivalence

$$\widehat{CF}(Y) \simeq \mathrm{Cone}\left(F_W \colon \widehat{CF}(Y_\lambda(K)) \to \widehat{CF}(Y_{\lambda+1}(K))\right). \tag{1.1}$$

(F_W denotes the "cobordism map", which we will discuss later).

Exercise 1.7. Recall that if $f\colon X \to Y$ is a chain map between chain complexes, then the mapping cone of f, denoted $\mathrm{Cone}(f)$, is defined to be the chain complex $X \oplus Y$, equipped with the differential

$$\begin{pmatrix} \partial_X & 0 \\ f & \partial_Y \end{pmatrix}.$$

Show that this is a differential, and that there is a short exact sequence

$$0 \to Y \to \mathrm{Cone}(f\colon X \to Y) \to X \to 0.$$

The snake lemma gives a long exact sequence on homology groups. Show that the connecting homomorphism is $H(f)$, the map on homology induced by f.

We now try to understand the Heegaard diagrams of the manifolds appearing in the statement about the exact triangles.

We pick a Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}_\infty)$ for Y. We assume that the Heegaard diagram is chosen so that K may be embedded on Σ, and there is a distinguished curve β_∞ which intersects K in a single point. Note that K may intersect $\boldsymbol{\alpha}$ many times, and we have no restriction here. The Heegaard surface Σ determines a framing for K, which we assume coincides with λ. Let β_λ be a curve parallel to K, and let $\beta_{\lambda+1}$ be obtained by twisting along a meridian of K, once (the sign of the twist is important). See Figure 1.1. We let $\boldsymbol{\beta}_\lambda$ and $\boldsymbol{\beta}_{\lambda+1}$ be the attaching curves obtained by replacing $\beta_\infty \in \boldsymbol{\beta}_\infty$ with β_λ and $\beta_{\lambda+1}$, respectively, but keeping the other curves.

We observe that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}_\lambda)$ represents $Y_\lambda(K)$ and $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}_{\lambda+1})$ represents $Y_{\lambda+1}(K)$.

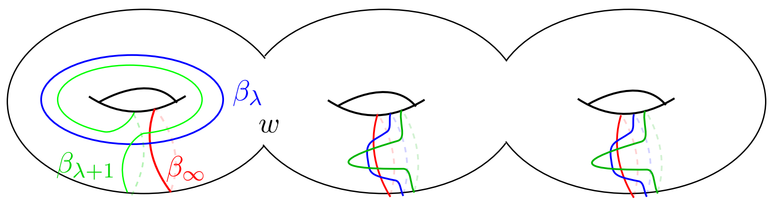

FIGURE 1.1. The attaching curves $\boldsymbol{\beta}_\infty$, $\boldsymbol{\beta}_\lambda$ and $\boldsymbol{\beta}_{\lambda+1}$.

Firstly, note that each of $\widehat{CF}(\Sigma, \boldsymbol{\beta}_\infty, \boldsymbol{\beta}_\lambda)$, $\widehat{CF}(\Sigma, \boldsymbol{\beta}_\lambda, \boldsymbol{\beta}_{\lambda+1})$, $\widehat{CF}(\Sigma, \boldsymbol{\beta}_{\lambda+1}, \boldsymbol{\beta}_\infty)$ is isomorphic to

$$\Lambda^{g-1},$$

as a graded group. (The exterior algebra on $g-1$ generators).

Exercise 1.8. Show that the differential vanishes.

In particular, we obtain canonical top degree intersection points $\Theta_{\infty,\lambda}$, $\Theta_{\lambda,\lambda+1}$ and $\Theta_{\lambda+1,\infty}$.

Using the above cycles, we can define a map

$$f_{\alpha,\beta_\lambda,\beta_{\lambda+1}}(-, \Theta_{\lambda,\lambda+1}) \colon \widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}_\lambda) \to \widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}_{\lambda+1}).$$

Lemma 1.9. *The map $f_{\alpha,\beta_\lambda,\beta_{\lambda+1}}(-, \Theta_{\lambda,\lambda+1})$ is a chain map.*

Proof. This follows from the A_∞-relations for holomorphic triangles, which imply

$$\partial f_{\alpha,\beta_\lambda,\beta_{\lambda+1}}(x, y) + f_{\alpha,\beta_\lambda,\beta_{\lambda+1}}(\partial x, y) + f_{\alpha,\beta_\lambda,\beta_{\lambda+1}}(x, \partial y) = 0$$

for any x and y. The main claim now follows from the fact that $\partial \Theta_{\lambda,\lambda+1} = 0$. □

Next, we define a map

$$\Phi \colon \widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}_\infty) \to \operatorname{Cone}(\widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}_\lambda) \to \widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}_{\lambda+1})),$$

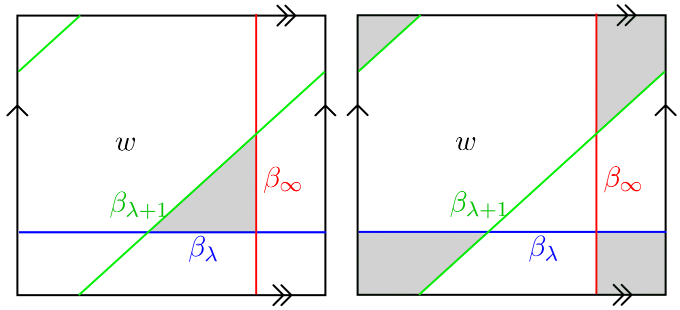

FIGURE 1.2. The figure $(\mathbb{T}^2, \beta_\infty, \beta_\lambda, \beta_{\lambda+1})$. Two triangle classes are shown.

via the following diagram:

$$\begin{tikzcd}
\widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}_\infty) \arrow[d, "f_{\alpha,\beta_\infty,\beta_\lambda}(-,\Theta)"] \arrow[drr, "f_{\alpha,\beta_\infty,\beta_\lambda,\beta_{\lambda+1}}(-,\Theta,\Theta)"] & & \\
\widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}_\lambda) \arrow[rr, "f_{\alpha,\beta_\lambda,\beta_{\lambda+1}}(-,\Theta)"] & & \widehat{CF}(\boldsymbol{\alpha}, \boldsymbol{\beta}_{\lambda+1})
\end{tikzcd}$$

Lemma 1.10. *The map Φ is a chain map.*

Proof. Using associativity of holomorphic quadrilaterals, this reduces to the claim that $f_{\boldsymbol{\beta}_\infty,\boldsymbol{\beta}_\lambda,\boldsymbol{\beta}_{\lambda+1}}(\Theta, \Theta) = 0$. The claim holds for arbitrary genus $g(\Sigma)$, but we are going to focus first on the case that $g = 1$. In this case, there are two triangles which make canceling contribution. See Figure 1.2. $\square$

By similar logic, there is a chain map Ψ in the opposite direction. We now consider the compositions $\Phi \circ \Psi$ and $\Psi \circ \Phi$. Since the polygon maps require attaching curves to be transverse, we write $\boldsymbol{\beta}'_\infty$, $\boldsymbol{\beta}'_\lambda$ and $\boldsymbol{\beta}'_{\lambda+1}$ for suitable small translates of these curves.

Lemma 1.11. *The compositions $\Phi \circ \Psi$ and $\Psi \circ \Phi$ are chain homotopic to the identity.*

Proof. We compute $\Psi \circ \Phi$ first. Associativity shows that the composition of the two maps is chain homotopic to

$$f_{\alpha,\beta_\infty,\beta'_\infty}(-, f_{\beta_\infty,\beta_\lambda,\beta_{\lambda+1},\beta'_\infty}(\Theta, \Theta, \Theta)).$$

We claim that $f_{\beta_\infty,\beta_\lambda,\beta_{\lambda+1},\beta'_\infty}(\Theta, \Theta, \Theta) = \Theta_{\beta_\infty,\beta'_\infty}$. This quadrilateral count is verified in Figure 1.3. Together with the subsequent lemma, the proof is complete. $\square$

Lemma 1.12. *If $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}')$ is a Heegaard triple where $\boldsymbol{\beta}'$ are "small-translates" of $\boldsymbol{\beta}$, then the map*

$$\mathbf{x} \mapsto f_{\alpha,\beta,\beta'}(\mathbf{x}, \Theta_{\beta,\beta'})$$

is a homotopy equivalence.

We will not give the proof, but we will give a moral reason for the proof. This is because given $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, there is a canonical "nearest-point" $\mathbf{x}_{np} \in \mathbb{T}_\alpha \cap \mathbb{T}_{\beta'}$. Furthermore, there are canonical "small-triangle" classes which clearly have a unique representative. See Figure 1.3. This argument maps it look like $f_{\alpha,\beta,\beta'}(\mathbf{x}, \Theta_{\beta,\beta'}) = \mathbf{x}_{np}$. *A-priori* there could be more triangles (e.g. with larger area on Σ), but it turns out one can show using Floer theoretic arguments that $f_{\alpha,\beta,\beta'}(-, \Theta_{\beta,\beta'})$ is chain homotopic to the "nearest-point map". Alternatively one can use a filtration argument using areas of triangles. We will not explore the details here. Compare Remark 2.11, below.

A very similar argument to the above shows that $\Phi \circ \Psi \simeq \mathrm{id}$.

Exercise 1.13. Verify that $\Phi \circ \Psi \simeq \mathrm{id}$.

After all of this, we are left with the following question:

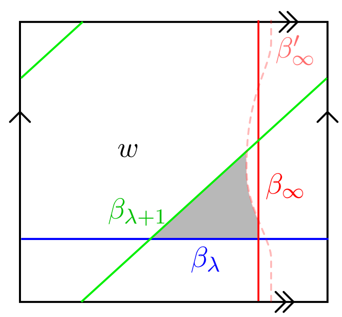

FIGURE 1.3. A rectangle on $(\mathbb{T}^2, \beta_\infty, \beta_\lambda, \beta_{\lambda+1}, \beta'_\infty)$.

Question 1.14. Why does this work? What is the most fundamental statement about the Lagrangians β_∞, β_λ and $\beta_{\lambda+1}$ which encodes the mapping cone formula and all of the model computations?

We will now state the answer, which we will unpack and prove in the next lecture:

Theorem 1.15. *There is a homotopy equivalence in the Fukaya category of* $\mathbb{T}^2$*:*

$$\beta_\infty \simeq \mathrm{Cone}(\theta\colon \beta_\lambda \to \beta_{\lambda+1}).$$

2. Lecture 2: Twisted complexes and the exact triangle

In this lecture, we are going to try to make sense of the following statement:

Theorem 2.1. *There is a homotopy equivalence in the Fukaya category of* $\mathbb{T}^2$*:*

$$\beta_\infty \simeq \mathrm{Cone}(\theta\colon \beta_\lambda \to \beta_{\lambda+1}).$$

2.1. **Fukaya category.** If (W, ω) is a symplectic manifold then there is an A_∞-category $\mathrm{Fuk}(W)$. (Here, we are assuming that suitably topological assumptions are satisfied to prevent pathological bubbling, but we will not worry about spelling details out here). We refer the reader to [Aur13] for a better introduction, and to [Sei08] for an excellent systematic treatment.

To establish the notation. A *Lagrangian* $L \subseteq W$ is a half-dimensional submanifold such that $\omega|_L = 0$. If L and L' are transverse Lagrangians, then

$$\operatorname{Hom}(L, L') = \operatorname{Hom}_{\mathrm{Fuk}}(L, L') := CF(L, L'),$$

which is the vector space (over $\mathbb{F} = \mathbb{Z}/2$) generated by intersection points. The spaces $\operatorname{Hom}(L, L')$ are themselves chain complexes, where the y-component of $\partial(x)$ is equal to the count of pseudo-holomorphic strips $u\colon \mathbb{D} \to W$ such that $u(-i) = x$, $u(i) = y$, and such that

$$u(S^1 \cap \{\Re(z) > 0\}) \subseteq L \quad \text{and} \quad u(S^1 \cap \{\Re(z) < 0\}) \subseteq L'.$$

Additionally we only want to count curves u which are part of a 1-dimensional moduli space. Such moduli spaces have a natural $\mathbb{R}$-action (by translation in the source), which we quotient by to get a number in $\mathbb{Z}/2$. The expected dimension of the moduli space is encoded by an algebro-topological quantity called the *Maslov index*, which we will not talk about here.

There is a natural "composition" map

$$\mu_2\colon \operatorname{Hom}(L', L'') \otimes \operatorname{Hom}(L, L') \to \operatorname{Hom}(L, L'')$$

which counts pseudo-holomorphic triangles in 0-dimensional moduli spaces. Here the z component of $\mu_2(x, y)$ is the count of triangles of index 0 (i.e. in 0-dimensional moduli spaces) with x, y and z along the boundary.

We should think of μ_2 as being parallel to "composition" in ordinary categories, such as the category of vector spaces, where we can compose two functions. For morphisms of vector spaces, we have strict associativity, i.e.

$$f \circ (g \circ h) = (f \circ g) \circ h,$$

however this fails for holomorphic polygons. Instead, we have associativity only up to chain homotopy. For example, we have

$$\mu_2(\mu_2(x \otimes y) \otimes z) + \mu_2(x \otimes \mu_2(y \otimes z)) + \mu_3(\partial(x \otimes y \otimes z)) + \partial\mu_3(x \otimes y \otimes z) = 0$$

where

$$\partial(x \otimes y \otimes z) = \partial(x) \otimes y \otimes z + x \otimes \partial(y) \otimes z + x \otimes y \otimes \partial(z).$$

Typically we write μ_1 instead of ∂, for aesthetic reasons.

In general, we have the A_∞ relations for any number of inputs, which read

$$\sum_{1 \le i < j \le n+1} \mu_{n-j+i}(x_1, \dots, x_{i-1}, \mu_{j-i+1}(x_i, \dots, x_j), x_{j+1}, \dots, x_n) = 0.$$

2.2. **Twisted complexes.** Our goal is to understand how to do homological algebra inside of the Fukaya category, i.e. we wish to define a *chain complex of Lagrangians*. More generally, if $\mathcal{C}$ is a category, we would naturally like to understand the following question:

Question 2.2. What is a chain complex in the category $\mathcal{C}$?

Remark 2.3. For the purposes of this note, we assume that we work over the field of 2 elements (to simplify signs).

If $\mathcal{C}$ is a category of R-modules for some ring R, we already knot the answer. It is a collection of spaces C_n and maps $\partial_n\colon C_n \to C_{n-1}$ such that $\partial_{n-1} \circ \partial_n = 0$. By considering $C = \bigoplus_n C_n$, we could instead just consider a pair (C, ∂) where $\partial\colon C \to C$ is an endormophism such that $\partial^2 = 0$.

Note that many constructions from homological algebra do not require that C actually be a vector space. An exception is taking homology $H_*(C)$ requires that we are working with vector spaces, $\ker / \operatorname{im}$. Note that "computing" homology could be rephrased by asking for which groups G is $H_*(C)$ isomorphic to. Equivalently, if we equip G with vanishing differential, we can rephrase this problem as asking for which chain complexes (C', ∂') is there an isomorphism $H_*(C', \partial') \cong H_*(C, \partial)$. This is an equivalence relation on chain complexes.

Note that there is a stronger equivalence relation between chain complexes called *homotopy equivalence*. Here, $C \simeq C'$ if we have chain maps $f\colon C \to C'$ and $g\colon C \to C'$ which are chain maps and $f \circ g = \mathrm{id} + \partial h + h\partial$ and $g \circ f = \mathrm{id} + \partial j + j\partial$. Note that $C \simeq C'$ implies that $H_*(C) \simeq H_*(C')$, so this can be thought of as a better behaved version of "computing homology". If we think about this, we observe that the fact that C and C' are themselves groups is not needed for this definition. Rather, the above statement is merely a consequence of the fact that $\operatorname{Hom}(C, C')$ (the set of linear maps from C to C') is itself a chain complex.

Exercise 2.4. If C and C' are chain complexes, then $\operatorname{Hom}(C, C')$ is a chain complex with differential $\partial(f) = \partial_{C'} \circ f + f \circ \partial_C$. (Note that when working over characteristic not equal to 2, there is a sign depending on the grading of f). Show that $H_*(\operatorname{Hom}(C, C'))$ is exactly the set of chain maps, modulo chain homotopy.

If $\mathcal{C}$ is a category, then we could define a *chain complex* in $\mathcal{C}$ to be a collection of objects C_n, together with morphisms ∂_n (in the category) such that $\partial_{n-1} \circ \partial_n = 0$. It turns out that this not quite general enough for our purposes. One issue is that in an A_∞-category like the Fukaya category, where we have morphisms, but not a natural notion of "compositions". Rather, we have "higher compositions" μ_j, and all of these must be taken into account.

Our final notion of a "chain complex in category $\mathcal{C}$" will be called a *twisted complex in* $\mathcal{C}$. We now embark on this notion. Suppose $\mathcal{C}$ is an A_∞-category. For convenience, we assume the morphism spaces are vector spaces over $\mathbb{F} = \mathbb{Z}/2$. The *additive enlargement* $\Sigma\mathcal{C}$ is as follows (see [Sei08, Sections 3k,l] for further details). Objects of $\Sigma\mathcal{C}$ consist of formal, finite collections of objects from $\mathcal{C}$ with formal grading shifts, and repeated elements allowed. More formally, one defines objects of the category to consist of collections $(X_i, V_i)_{i \in I}$ such that I is a finite index set, each X_i is an object of $\mathcal{C}$ and each V_i is a finite dimensional, graded vector space. If X and Y are objects of $\Sigma\mathcal{C}$, then $\operatorname{Hom}(X, Y)$ is defined to be "matrices" of morphisms between the objects and vector spaces of X and Y. I.e. if $X = (X_i, V_i)_{i \in I}$ and $Y = (Y_j, W_j)_{j \in J}$ are objects of $\Sigma\mathcal{C}$, then $\operatorname{Hom}(X, Y)$ is defined to be the direct sum over $(i, j) \in I \times J$ of $\operatorname{Hom}(V_i, W_j) \otimes \operatorname{Hom}_{\mathcal{A}}(X_i, Y_j)$. It is straightforward to verify that $\Sigma\mathcal{C}$ is naturally an A_∞-category.

Definition 2.5. A *twisted complex* in $\mathcal{C}$ consists of an object X of $\Sigma\mathcal{C}$, together with an endomorphism $\delta_X \in \operatorname{Hom}(X, X)$ of degree -1, such that

$$\sum_{n \geq 1} \mu_n^{\Sigma\mathcal{C}}(\delta_X, \dots, \delta_X) = 0. \tag{2.1}$$

Note that some assumption is necessary to ensure that the above sum is finite, and this is usually in the definition of the category. Different assumptions are useful in different contexts, so we will not make this formal. Note that if $\mu_n^{\Sigma\mathcal{C}} = 0$ for sufficiently large n, then finiteness is achieved.

Twisted complex $\mathrm{Tw}(\mathcal{C})$ naturally form an A_∞-category. Morphisms are the same as in $\Sigma\mathcal{C}$. Given a composable sequence of morphisms

$$X_0 \xrightarrow{f_{0,1}} \cdots \xrightarrow{f_{n-1,n}} X_n,$$

one defines

$$\begin{aligned}&\mu_n^{\mathrm{Tw}}(f_{0,1},\dots,f_{n-1,n})\\ &= \sum_{i_0,\dots,i_n\geq 0} \mu^{\Sigma\mathcal{C}}_{n+i_0+\cdots+i_n}(\overbrace{\delta_{X_0},\dots,\delta_{X_0}}^{i_0}, f_{0,1}, \overbrace{\delta_{X_1},\dots,\delta_{X_1}}^{i_1},\dots, f_{n-1,n}, \overbrace{\delta_{X_n},\dots,\delta_{X_n}}^{i_n})\end{aligned}$$

Exercise 2.6. Prove that if X_0 and X_1 are twisted complexes, then $\mathrm{Hom}(X_0, X_1)$ is a chain complex. More generally, verify that if $\mathcal{C}$ is an A_∞ category, then $\mathrm{Tw}(\mathcal{C})$ is an A_∞-category. If $\mathcal{C}$ is an A_∞-category which has $\mu_j = 0$ for $j > 2$ (i.e. each $\mathrm{Hom}(X, Y)$ is a chain complex, composition is strictly associative and satisfies the Leibniz rule), then show that $\mathrm{Tw}(\mathcal{C})$ has the same property.

If $f\colon X \to Y$ is a morphism of twisted complexes which satisfies $\mu_1(f) = 0$, then we may naturally construct the mapping cone $\mathrm{Cone}(f)$. The underlying space of the cone is the union of the elements of X, and δ_X is obtained by the union of δ_X, δ_Y and f (each viewed as a set of component maps ranging over the components X_i and Y_j of X and Y).

Exercise 2.7. Show that if $f\colon X \to Y$ is a map of twisted complexes which satisfies $\mu_1(f) = 0$, then $\mathrm{Cone}(f)$ is a twisted complex (in particular, Equation (2.1) is satisfied).

Twisted complexes appear very frequently in homological algebra. Here are a few important examples:

(1) Chain complexes are twisted complexes of modules.
(2) Twisted complexes over an algebra (see next section), naturally encode projective dg-modules over that algebra.
(3) Twisted complexes in the Fukaya category are the natural home of our surgery exact triangles.
(4) Twisted complexes of vector bundles are an important part of K-theory.

2.2.1. *Twisted complexes over algebras.* We now discuss twisted complexes over algebra. These encode familiar notions from homological algebra, and should be thought of as "projective chain complexes".

By an *idempotent ring* $\mathbf{i}$ we mean a finite direct product of the rings $\mathbb{F}$. Suppose that $\mathcal{A}$ is an A_∞-algebra over an idempotent ring $\mathbf{i}$, then we may view $\mathcal{A}$ as an A_∞ category whose objects are idempotents $i \in \mathbf{i}$, and such that $\mathrm{Hom}(i,j) = i \cdot \mathcal{A} \cdot j$. The fact that $\mathcal{A}$ is an A_∞-algebra is equivalent to the A_∞-category relations.

It is helpful to discuss the category of twisted complexes in more detail. These take the form of a formal collection of idempotents (with grading shifts). Equivalently, we can think of such a collection as a finitely generated right $\mathbf{i}$-module X. We have a collection of morphisms $f_{i,j}$, ranging over the generators. We can package this as a map

$$\delta^1 \colon X \to \mathcal{A} \otimes_{\mathbf{i}} X.$$

It is helpful to define a map $\delta^n \colon X \to \mathcal{A}^{\otimes n} \otimes X$ as the n-fold iterate of δ^1. Then the twisted complex relation is equivalent to

$$\sum_{n \geq 1} (\mu_n \otimes \mathrm{id}_X) \circ \delta^n = 0.$$

Schematically,

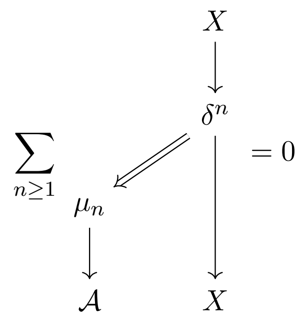

In all cases, there is a well-defined differential on $\mathcal{A} \otimes_{\mathbf{i}} X$, given by

$$\partial(a \otimes x) = \sum_{n \geq 0} (\mu_{n+1}(a, -) \otimes \mathrm{id}_X)(\delta^n(x)),$$

(where $\delta^0 = \mathrm{id}$), or schematically

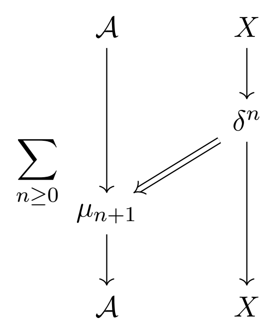

Since $\mathcal{A} \otimes_{\mathbf{i}} X$ embeds as a direct sum (over $\mathbf{i}$) of copies of $\mathcal{A}$, so we should think of a type-D structure as a projective $\mathcal{A}$-module.

It is helpful to think about this when $\mathcal{A}$ is an algebra, resp. *dg*-algebra:

(1) If $\mathcal{A}$ is an algebra, then the type-D relations are just

$$(\mu_2 \otimes \mathrm{id}) \circ (\mathrm{id}_{\mathcal{A}} \otimes \delta^1) \circ \delta^1 = 0.$$

(2) If $\mathcal{A}$ is a dg-algebra, then

$$(\mu_1 \otimes \mathrm{id}_X) \circ \delta + (\mu_2 \otimes \mathrm{id}) \circ (\mathrm{id}_{\mathcal{A}} \otimes \delta^1) \circ \delta^1 = 0.$$

Remark 2.8. If $\mathcal{A}$ is an ordinary algebra over the idempotent ring $\mathbb{F}$ (i.e. one idempotent) then a type-D module over A is the same a free chain complex over $\mathcal{A}$.

Exercise 2.9. Show that when $\mathcal{A}$ is an algebra (resp. *dg*-algebra) then $Y = \mathcal{A} \otimes X$ is naturally a differential graded module, i.e.

$$\partial(a \cdot y) = \partial(a) \cdot y + a \cdot \partial(y).$$

If $\mathcal{A}$ is an A_∞-algebra, verify that $\mathcal{A} \otimes X$ is naturally a left A_∞-module over $\mathcal{A}$.

Remark 2.10. Type-D structures over an algebra are morally the same as projective modules, and for practical purposes they are interchangeable (e.g. for computing derived functors and tensor products). Note however there are some differences, and generally type-D structures have some practical advantages. For example, over a polynomial ring, a famous and exceptionally hard theorem of Quillen-Suslin says that all projective modules over $\mathbb{F}[X_1, \dots, X_n]$ are free. Type-D structures over $\mathbb{F}[X_1, \dots, X_n]$ are trivially free.

2.3. **Twisted complexes in the Fukaya category.** Next, we investigate twisted complexes in the Fukaya category. Instead of idempotents, the objects are now Lagrangians L. Note that some requirements are necessary to form twisted complexes, since $CF(L, L')$ usually requires some version of *admissibility* to be defined.

We illustrate some examples and explain a few subtleties:

(1) If L and L' are Lagrangians, then a diagram

$$L \xrightarrow{\mathbf{x}} L'$$

is a twisted complex iff $\mathbf{x}$ is a cycle.

(2) If L_0, L_1 and L_2 are Lagrangians, then a diagram

$$\begin{tikzcd} L_0 \arrow[r, "\mathbf{x}"] \arrow[dr, "\mathbf{z}"] & L_1 \arrow[d, "\mathbf{y}"] \\ & L_2 \end{tikzcd}$$

is a twisted complex iff

$$\mu_2(\mathbf{y}, \mathbf{x}) + \mu_1(\mathbf{z}) = 0.$$

I.e. the composition $\mu_2(\mathbf{y}, \mathbf{x})$ is a null-homotopic, with $\mathbf{z}$ being the null-homotopy.

Remark 2.11. You may observe that the Fukaya category naturally lacks identity elements in $CF(L, L)$, since we always assume that L and L' are transverse for $CF(L, L')$ to be defined. (And L is never transverse to itself). This turns out to be less of an issue than it seems. It turns out that the correct perspective is that instead of units (or "strict units") there are "homotopy units". We are mostly interested in the case that L is a 1-sphere (or torus), in which case we let L' be a small translation which intersects L transversely in two points and we let $\theta_1 \in CF(L, L')$ be the top graded intersection point. The element θ_1 is a homotopy identity, in the sense that $\mu_2(\theta_1, -)$ is always a homotopy equivalence. It turns out that from a categorical perspective, "homotopy unital" categories can always be deformed to "strictly unital" categories (see [Sei08, Section 2a] for precise details), and so we are not committing a crime by just pretending that θ_1 is a strict unit.

2.4. **Exercises.** In the following, we work over characteristic 2 and ignore grading shifts.

Exercise 2.12. Given a short exact sequence of chain complexes

$$0 \to A \to B \to C \to 0,$$

show that the following maps are quasi-isomorphism:

(1) $A \to \mathrm{Cone}(B \to C)$.
(2) $\mathrm{Cone}(A \to B) \to C$.
(3) If C is a projective R-module (e.g. a vector space over a field), show that all of the above are homotopy equivalences.
(4) Assuming C is projective, show that there is a chain map $f\colon C \to A$ such that $B \simeq \mathrm{Cone}(f\colon C \to A)$.

2.5. **Twisted complexes and the exact triangle.** Recall that our goal is to show that

$$\gamma_2 \simeq \mathrm{Cone}(\theta\colon \gamma_0 \to \gamma_1).$$

For this, we define a morphism

$$\Phi\colon \gamma_2 \to \mathrm{Cone}(\theta\colon \gamma_0 \to \gamma_1).$$

This morphism is the unique intersection point (which we denote by ϕ) between γ_2 and γ_0, i.e. the following diagram:

$$\begin{tikzcd} \gamma_2 \arrow[d, "\phi"] & \\ \gamma_0 \arrow[r, "\theta"] & \gamma_1 \end{tikzcd}$$

Lemma 2.13. *The map Φ is a cycle (i.e. chain map).*

Proof. It suffices to check that

$$\partial(\phi) = 0 \quad \text{and} \quad \mu_2(\phi, \theta) = 0.$$

To check that $\partial(\phi) = 0$ is easy. There are no disks. The computation that $\mu_2(\phi, \theta) = 0$ is more subtle is exactly the computation shown in Figure 1.2. □

Remark 2.14. When working over CF^-, the exact triangle only holds with coefficients in $\mathbb{F}[\![U]\!]$, and not $\mathbb{F}[U]$. We can see above the importance of working with power series because we indeed do have infinitely many holomorphic triangles which contribute, and the sums only converge if we work in the power series ring.

Next, we can define a morphism in $\mathrm{Tw}(\mathrm{Fuk}(\mathbb{T}^2))$ backwards

$$\Psi\colon\ \mathrm{Cone}(\theta\colon \gamma_0 \to \gamma_1) \to \gamma_2.$$

It is important to actually trasnlate the second copy of γ_2 very slightly so we consider Ψ as a map from $\mathrm{Cone}(\theta\colon \gamma_0 \to \gamma_1) \to \gamma_2'$, as before. We define ψ to be the unique intersection point of γ_2 and γ_0'.

We now consider the composition $\Psi\circ\Phi := \mu_2(\Psi,\Phi)$. We claim that this is homotopic to the morphism θ^+.

$$\begin{tikzcd}
\gamma_0 \arrow[d, "\phi"] & \\
\gamma_1 \arrow[r, "\theta"] & \gamma_2 \arrow[d, "\psi"] \\
& \gamma_0'
\end{tikzcd}$$

We can compute the composition $\mu_2(\Psi,\Phi)$. By definition, this is $\mu_3(\phi,\theta,\psi)$. The model computation from Figure 1.3 now identifies this $\mu_3(\phi,\theta,\psi)=\theta^+$, the top graded intersection point between γ_0 and γ_0'.

Next, we consider the composition $\mu_2(\Phi,\Psi)$. In this case, the diagrams are a bit more complicated

$$\begin{tikzcd}
\gamma_1 \arrow[r, "\theta"] & \gamma_2 \arrow[d, "\psi"] & \\
& \gamma_0' \arrow[d, "\phi"] & \\
& \gamma_1' \arrow[r, "\theta"] & \gamma_2'.
\end{tikzcd}$$

Exercise 2.15. Verify that the composition $\mu_2(\Phi,\Psi)$ is the diagram (viewed as a morphism from the top line to the bottom line)

$$\begin{tikzcd}
\gamma_1 \arrow[r, "\theta"] \arrow[d, "\theta^+"] & \gamma_2 \arrow[d, "\theta^+"] \\
\gamma_1' \arrow[r, "\theta"] & \gamma_2'
\end{tikzcd}$$

We now consider these claims over CF^-, which we recall is a finitely generated free complex over $\mathbb{F}[\![U]\!]$. In this version, we allow curves to pass over the basepoint w, and we have an additional factor of $U^{n_w(\phi)}$. The model computations can still be pursued in this theory. In this case, there are actually quite a lot more holomorphic curve classes (in fact infinitely many). We can still count them without too much trouble. To count them, we note that a holomorphic triangle is a map $u\colon \Delta\to\mathbb{T}^2$, so in this case we can lift triangles to the universal cover $\widetilde{u}\colon \Delta\to\mathbb{R}^2$, which makes things easier to visualize. Figure 1.2 generalizes, and we see that there are triangles u_s^+ and u_s^-, with $s\in\{0,1,2,\dots,\}$. These triangles have $n_w(u_s^+)=n_w(u_s^-)=s(s-1)/2$ and hence u_s^+ and u_s^- are both weighted by $U^{s(s-1)/2}$ and hence cancel modulo 2.

Note that it is clear now why we need the power series ring instead of the polynomial ring (there are infinitely many curves to count). It is possible (though we will not investigate it) to show that there are only finitely many curves with a given value of $n_w(\phi)$, so the maps are still sensible.

Exercise 2.16. Extend all of the model computation in this section for CF^-. You should be able to completely answer things for the $\mathbb{T}^2$ portion of the computation. The full genus g computations require some extra steps, which we will not embark on understanding in these notes. The answer that one should arrive at is that $\mu_2(\Phi,\Psi)$ and $\mu_2(\Psi,\Phi)$ are

$$\sum_{s\geq 0} U^{s(s-1)/2}\cdot \theta^+.$$

Observe that that the coefficient is a unit in $\mathbb{F}[\![U]\!]$.

3. Lecture 3: L-space knots and links

In this lecture, we consider the question of computing and understanding the knot Floer complexes for knots in S^3. We can view these as complexes $\mathcal{CFK}(K)$ which are free and finitely generated over $\mathbb{F}[\mathscr{U}, \mathscr{V}]$.

3.1. **Lens spaces.** A *lens space* is one of the simplest 3-manifolds. There are a variety of ways of defining a lens space, but the simplest is that it is a 3-manifold with a genus 1 Heegaard splitting. (Sometimes $S^1 \times S^2$ is not counted as a lens space).

Let p and q be coprime. We pick an oriented basis (μ, λ) of $\mathbb{T}^2$. Then $L(p,q)$ is the 3-manifold with a Heegaard splitting $(\mathbb{T}^2, \alpha, \beta)$ where $\alpha = \mu$ and $\beta = p\lambda + q\mu$. See below:

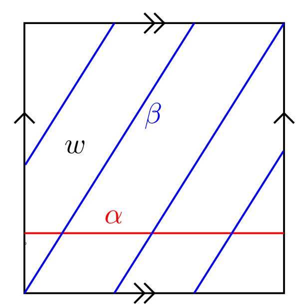

Figure 3.1. A Heegaard diagram for the lens space $L(3,2)$.

Lemma 3.1. *$\widehat{HF}(L(p,q)) \cong \oplus_p \mathbb{F}$ and $HF^-(L(p,q)) \cong \oplus_p \mathbb{F}[U]$.*

Proof. The above diagram has p intersection points. The differential vanishes since there are exist no classes of disks which connect different intersection points (let alone any holomorphic curves of index 1). $\square$

Remark 3.2. The different intersection points each represent a different Spin^c structure; see the next section.

Exercise 3.3. Show that pq surgery on the torus knot $T(p,q)$ is $L(p,q)\#L(q,p)$.

The above exercise is somewhat hard, but more details may be found in the work of Moser [Mos71, Proposition 4].

3.2. Spin^c **structures.** We will only very briefly discuss Spin^c structures. Let Y be a 3-manifold. Then TY, the tangent bundle, determines a homotopy class of maps

$$Y \to BO(3)$$

where $BO(3)$ is the classifying space of the orthogonal group. (The definition of classifying spaces is beyond the scope of these notes). An orientation of Y is equivalent to a choice of lift of the above map

$$\begin{tikzcd} & BSO(3) \arrow[d] \\ Y \arrow[ur, dashed] \arrow[r] & BO(3) \end{tikzcd}.$$

It turns out there is another Lie group which is important for our purposes, called $\mathrm{Spin}^c(3)$. (The definition is not particularly important for our notes). The important thing is that there is a short exact sequence of Lie groups

$$U(1) \to \mathrm{Spin}^c(3) \to SO(3),$$

which in turn induce maps on classifying spaces. A Spin^c-*structure* on Y is a homotopy class of lift

$$\begin{tikzcd} & B\,\mathrm{Spin}^c(3) \arrow[d] \\ Y \arrow[r] \arrow[ur, dashed] & BSO(3) \end{tikzcd}$$

The ambiguity of the choice of lift turns out to lie in $[Y, BU(1)] = [Y, \mathbb{CP}^\infty] \cong H^2(Y)$.

It turns out that Heegaard Floer homology naturally decomposes over Spin^c structures on Y, and we write

$$\widehat{HF}(Y) \cong \bigoplus_{\mathfrak{s} \in \mathrm{Spin}^c(Y)} \widehat{HF}(Y, \mathfrak{s}).$$

Remark 3.4. Ozsváth and Szabó prove that if Y is a rational homology 3-sphere, then $\widehat{HF}(Y, \mathfrak{s}) \neq 0$ for all $\mathfrak{s} \in \mathrm{Spin}^c(Y)$. Similarly, Ozsváth and Szabó prove that $HF^-(Y, \mathfrak{s})$ is always isomorphic to the direct sum of one copy of $\mathbb{F}[U]$ together with some number of copies of $\mathbb{F}[U]/U^n$, for various n.

3.3. **L-spaces and L-space knots.** We focus on L-space knots. Recall that a rational homology 3-sphere Y is an L-space if $\widehat{HF}(Y) \cong \bigoplus_{\mathrm{Spin}^c(Y)} \mathbb{F}$. Equivalently,

$$HF^-(Y) \cong \bigoplus_{\mathrm{Spin}^c(Y)} \mathbb{F}[U].$$

Remark 3.5. In general, if Y is a rational homology 3-sphere, then $\widehat{HF}(Y)$ has rank at least $|\,\mathrm{Spin}^c(Y)| = |H_1(Y)|$ (the number of elements in $H_1(Y)$).

Exercise 3.6. Using the structural description in Remark 3.4, prove that the above two notions of L-space are equivalent by examining the short exact sequence

$$0 \to CF^-(Y) \xrightarrow{U} CF^-(Y) \to \widehat{CF}(Y) \to 0$$

and the induced long exact sequence.

Definition 3.7. We say that $K \subseteq S^3$ is an *L-space knot* if there is an $n \in \mathbb{N}$ such that $S^3_n(K)$ is an L-space.

Exercise 3.8. Show (using the surgery exact triangle) that if $n > 0$, then $S^3_n(K)$ is an L-space, then $S^3_{n+1}(K)$ is an L-space. If $n < 0$, show that if $S^3_n(K)$ is an L-space then $S^3_{n-1}(K)$ is also an L-space.

Remark 3.9. All torus knots $T_{p,q}$ are L-spaces knots, since $pq \pm 1$ surgery on $T_{p,q}$ gives $L(pq \pm 1, q^2)$. Also, pq surgery on $T(p,q)$ is $L(p,q) \# L(q,p)$. More generally, Hedden has shown that all algebraic knots are L-space knots [Hed09].

We are interested in *staircase complexes*, which are free complexes over the 2-variable polynomial ring $\mathbb{F}[\mathscr{U},\mathscr{V}]$ of the following form:

$$\begin{tikzcd}[column sep=small]
& x_1 \arrow[dl, "\mathscr{U}^{a_1}"'] \arrow[dr, "\mathscr{V}^{b_1}"] & & x_3 \arrow[dl, "\mathscr{U}^{a_3}"'] \arrow[dr, "\mathscr{V}^{b_3}"] & & \cdots & & x_{2n-3} \arrow[dl, "\mathscr{U}^{a_{2n-3}}"'] \arrow[dr, "\mathscr{V}^{b_{2n-3}}"] & & x_{2n-1} \arrow[dl, "\mathscr{U}^{a_{2n-1}}"'] \arrow[dr, "\mathscr{V}^{b_{2n-1}}"] & \\
\mathbf{x}_0 & & \mathbf{x}_2 & & \cdots & \cdots & \cdots & & \mathbf{x}_{2n-2} & & \mathbf{x}_{2n}
\end{tikzcd}$$

In this lecture, we will prove the following result of Ozsváth and Szabó:

Theorem 3.10 ([OS05]). *The knot Floer complex of an L-space knot is a staircase.*

Note that for an L-space knot, no two generators lie in the same Alexander and Maslov gradings, and hence we can easily read off the complex from the Alexander polynomial.

Remark 3.11. Note that this gives a very strong restriction on which knots have lens space surgeries. For example it says that if K is a knot with a lens space surgery, then all coefficients in the Alexander polynomial $\Delta_K(t)$ are ± 1.

Note that the Alexander polynomial of $T(p,q)$ is

$$\Delta_{T(p,q)}(t) \doteq \frac{(t^{pq}-1)(t-1)}{(t^p-1)(t^q-1)}.$$

3.4. **More on L-space knots.** We now return to study L-space knots. There is an extra ingredient in Ozsváth and Szabó's proof, and that's a *large surgery formula.* They prove that $\mathcal{CFK}(K)$ naturally computes $HF^-(S^3_N(K))$ for N suitably large. To state their theorem, we need to consider gradings. Knot Floer homology naturally has two Maslov gradings, gr_w and gr_z. Here,

$$(\mathrm{gr}_w,\mathrm{gr}_z)(\mathscr{U}) = (-2,0) \quad \text{and} \quad (\mathrm{gr}_w,\mathrm{gr}_z)(\mathscr{V}) = (0,-2).$$

It is natural to consider a third grading (linearly dependent on the above)

$$A = \tfrac{1}{2}(\mathrm{gr}_w - \mathrm{gr}_z).$$

Inside of $\mathcal{CFK}(K)$, we consider the subspace in Alexander grading $s \in \mathbb{Z}$, denoted $A_s(K)$. Note that $A_s(K)$ is not preserved by either $\mathscr{U}$ or $\mathscr{V}$, since they have non-zero Alexander grading. However, it is preserved by the product

$$U = \mathscr{U}\mathscr{V}.$$

Note that we can therefore decompose $\mathcal{CFK}(K)$ as a direct sum (over $\mathbb{F}[U]$) of all of the $A_s(K)$:

$$\mathcal{CFK}(K) = \bigoplus_{s\in\mathbb{Z}} A_s(K).$$

Theorem 3.12 (Ozsváth Szabó). *If $K \subseteq S^3$ is a knot and $N \gg 0$, then*

$$H_*(A_s(K)) \cong HF^-(S^3_N(K),[s])$$

for all $s \in [-N/2, N/2]$.

See [OS08, Theorem 2.3].

In the above the number s determines a particular Spin^c structure $[s]$ on $S^3_N(K)$. Also, if you've forgotten, Heegaard Floer decomposes over Spin^c structures:

$$HF^-(Y) \cong \bigoplus_{\mathfrak{s}\in \mathrm{Spin}^c(Y)} HF^-(Y,\mathfrak{s}).$$

Note also that there are some additional restrictions on the knot Floer homology groups. There is a map

$$\mathcal{CFK}(K) \to \mathcal{CFK}(K)/(\mathscr{U}-1) \cong CF^-(S^3),$$

gotten by forgetting the w basepoint. This map is easily seen to send $\mathbb{F}[U]$-non-torsion elements to $\mathbb{F}[U]$-non-torsion elements, and also map the gr_z-grading with the Maslov grading on $HF^-(S^3)$. There is a similar map

$$\mathcal{CFK}(K) \to \mathcal{CFK}(K)/(\mathscr{V}-1) \cong CF^-(S^3).$$

By considering the two quotient maps, we obtain the following structural result:

Corollary 3.13. *If K is an L-space knot, then $H(\mathcal{CFK}(K))$ has homology supported only in $(\mathrm{gr}_w, \mathrm{gr}_z)$-bigradings (i,j) where $i,j \le 0$. Furthermore, it has $\mathbb{F}$-rank at most one in each grading.*

Definition 3.14. A *monomial ideal I* in $\mathbb{F}[\mathscr{U},\mathscr{V}]$ is an ideal spanned by monomials $\mathscr{U}^i\mathscr{V}^j$ ranging over various $i,j\ge 0$.

Corollary 3.15. *The homology $H(\mathcal{CFK}(K))$ is an monomial ideal in $\mathbb{F}[\mathscr{U},\mathscr{V}]$.*

Exercise 3.16. Show that every monomial ideal $I \subseteq \mathbb{F}[\mathscr{U},\mathscr{V}]$ has a free resolution consisting of two "steps". I.e. that there are free, finitely generated $\mathbb{F}[\mathscr{U},\mathscr{V}]$-modules C_1 and C_0 and an $\mathbb{F}[\mathscr{U},\mathscr{V}]$-linear map $f\colon C_1 \to C_0$ such that the below sequence is exact

$$0 \to C_1 \xrightarrow{f} C_0 \to I \to 0.$$

In Figure 3.2, we have an example of both $\mathcal{CFK}(K)$ and $H_*(\mathcal{CFK}(K))$.

We now sketch a proof of the fact that L-space knots have knot Floer complexes which are staircases. Our proof will use the following fact, which we will not prove:

Lemma 3.17. *If X and Y are finitely generated, free chain complexes over $\mathbb{F}[\mathscr{U},\mathscr{V}]$ and there is a quasi-isomorphism $f\colon X\to Y$, then X and Y are homotopy equivalent.*

Remark 3.18. Proofs of results similar to the above can be found in [Wei94, Section 10.4].

Exercise 3.19. Let $A = \mathbb{F}[X] \xrightarrow{X^n} \mathbb{F}[X]$ (where $\mathbb{F}$ is the field of 2-elements). There is a chain map $A \to \mathbb{F}[X]/X^n$ which is projection onto the second summand. Verify that this is a quasi-isomorphism and admits no homotopy inverse. If you are motivated, describe a homotopy inverse as an A_∞-module over the ring $\mathbb{F}[X]$.

We now give a useful algebraic lemma:

Lemma 3.20. *Let (X,∂_X) be a free chain complex over $\mathbb{F}[\mathscr{U},\mathscr{V}]$ whose homology $H(X)$ admits a 2-step free resolution. I.e. such that we have an exact sequence*

$$0 \to C_0 \xrightarrow{f} C_1 \to H(X) \to 0,$$

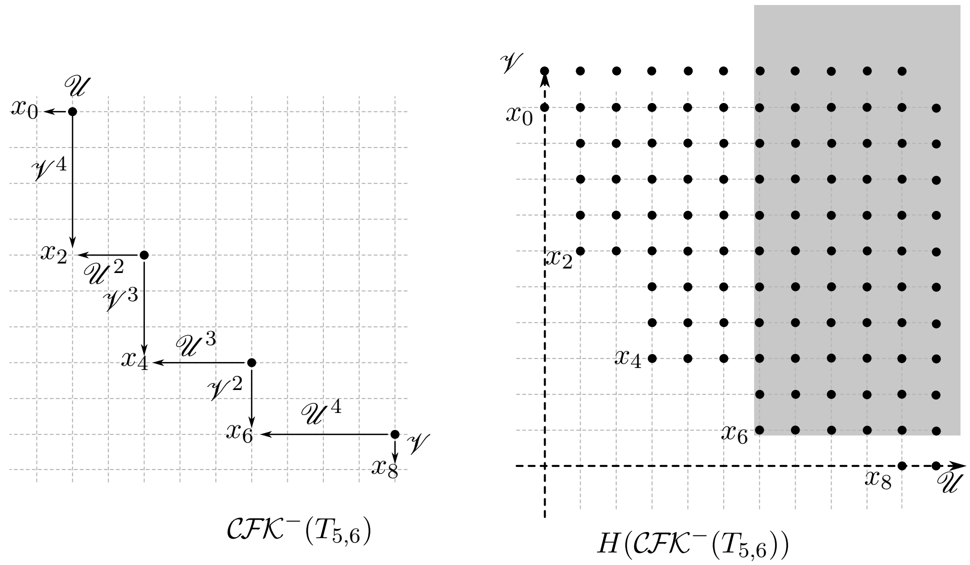

FIGURE 3.2. The complex $\mathcal{CFK}(T_{5,6})$ and its homology $H(\mathcal{CFK}(T_{5,6}))$, viewed as a subspace of $\mathbb{F}[\mathscr{U},\mathscr{V}]$. The Alexander polynomial is $\Delta_{T_{5,6}}(t) = t^{10} - t^9 + t^5 - t^3 + 1 - t^{-3} + t^{-5} - t^{-9} + t^{-10}$.

where C_i are free $\mathbb{F}[\mathscr{U},\mathscr{V}]$-modules (with trivial differential). Then X is quasi-isomorphic to $\mathcal{C} := \mathrm{Cone}(f\colon C_0 \to C_1)$.

Proof. We will define a quasi-isomorphism $\mathcal{C} \to X$, as follows. Let $B \subseteq Z \subseteq X$ denote the image and kernel of ∂_X, respectively. We pick an isomoprhism from $C_0/\operatorname{im}(C_1)$ to $H(X)$, and consider the diagram below:

$$\begin{tikzcd} 0 \arrow[r] & C_1 \arrow[r, "f"] \arrow[d, dashed, "\phi_1"] & C_0 \arrow[r] \arrow[d, dashed, "\phi_0"] & C_0/C_1 \arrow[r] \arrow[d, "\phi"] & 0 \\ & X \arrow[r, "\partial_X"] & Z \arrow[r] & H(X) \arrow[r] & 0. \end{tikzcd}$$

The rows are exact. Since $Z \to H(X)$ is surjective and C_0 is free, we can pick a $\phi_0\colon C_0 \to Z$ making the right square in the diagram commute. Next, we observe that $\phi_0 \circ f$ projects trivially into $H(X)$, by commutativity, so we can factor $f \circ \phi_0$ into $\operatorname{im}\partial_X = B$. Since ∂_X surjects as a map from X to B, and C_1 is free, we can pick a ϕ_1 making the diagram commute.

We may now define a chain map from $\mathcal{C}$ to X, given by (ϕ_1, ϕ_0). I.e. the map sends C_0 to X via ϕ_0, and C_1 to X via ϕ_1. It is straightforward to see that this map is a chain map. Furthermore, it induces the map ϕ on homology, so is a quasi-isomorphism. □

3.5. **Links.** One may also ask about L-space links instead of L-space knots. It turns out the situation is comparably more complicated. It is not the case if a link L has a single L-space surgery, then all sufficiently positive or sufficiently negative surgeries will be L-space links. Instead, the most natural condition is the following:

Definition 3.21. A link L is an *L-space link* if all sufficiently positive surgeries (i.e. all components are given sufficiently positive framing) on L are L-spaces.

Gorsky and Némethi have proven that all algebraic links are L-space links [GN16], and hence the above family is important.

There is a Heegaard link complex $\mathcal{CFL}(L)$, parallel to the knot Floer complexes. This complex has a $\mathscr{U}_i$ and $\mathscr{V}_i$ for each link component, and is a finitely generated free complex over the ring $\mathbb{F}[\mathscr{U}_1, \mathscr{V}_1, \dots, \mathscr{U}_n, \mathscr{V}_n]$ where $n = |L|$.

There is a large surgeries theorem parallel for links as well, and a simple consequence is the following:

Lemma 3.22. *A link L is an L-space link if and only if the homology $H(\mathcal{CFL}(L))$ is a torsion free $\mathbb{F}[U]$-module. (Here, U denotes any of $\mathscr{U}_i\mathscr{V}_i$; it turns out $\mathscr{U}_i\mathscr{V}_i$ and $\mathscr{U}_j\mathscr{V}_j$ are chain homotopic as endomorphisms so induce the same action on homology).*

The algebraic setting is comparably harder, and our previous argument does not work in this setting. Nonetheless, it can be proven [BLZ22] that the link Floer complexes $\mathcal{CFL}(L)$ of algebraic links are free-resolutions of their homology $H(\mathcal{CFL}(L))$, giving a parallel result to Ozsváth and Szabó's result for L-space knots.

3.6. **Exercises.**

(1) Consider the following two free chain complexes over $\mathbb{F}[\mathscr{U}, \mathscr{V}]$:

$$X = \bullet \oplus \begin{tikzcd} \bullet \arrow[r, "\mathscr{U}"] \arrow[d, "\mathscr{V}"] & \bullet \arrow[d, "\mathscr{V}"] \\ \bullet \arrow[r, "\mathscr{U}"] & \bullet \end{tikzcd} \qquad \text{and} \qquad Y = \begin{tikzcd} & \bullet \arrow[d, "\mathscr{V}"] \\ \bullet \arrow[r, "\mathscr{U}"] & \bullet \end{tikzcd}$$

Verify that complexes of X and Y have isomorphic homology as modules over $\mathbb{F}[\mathscr{U}, \mathscr{V}]$ (here we are ignoring gradings; but the claim holds for some choice of relative grading), but that the underlying complexes X and Y are not homotopy equivalent. If you are ambitious, you can also compute the induced A_∞-module structure on the two homology groups.

(2) Compute the induced A_∞-module structure on the homology of X and Y.

4. Lecture 5: Plumbed 3-manifolds

We investigate the Heegaard Floer homologies of some plumbed 3-manifolds. The main references of [OS03] [Ném05] [Ném08].

4.1. **Plumbed 3-manifold and 4-manifolds.**

Let G be a tree with integer weights at the vertices. There is an associated 4-manifold $X(G)$ with boundary $Y(G)$ (a 3-manifold). We call $Y(G)$ a *plumbed 3-manifold*. We describe the construction below.

If v is a vertex, we write $w(v)$ for the weight. For each vertex v we take a disk bundle over S^2 with Euler number $w(v)$. (Note that the boundary is a lens space $L(\pm v, 1)$, with $\pm$ determined by orientation conventions). A disk bundle of Euler number n is a fiber bundle

$$D^2 \to W(n) \to S^2.$$

Remark 4.1. The Euler number is the self-intersection number of the 0-section, which is a copy of S^2.

Given a small disk $D \subseteq S^2$, we may trivialize this bundle, and obtain a copy of $D^2 \times D^2$ inside of $W(n)$. The first factor of D^2 is the *base*, and the second copy is the *fiber*.

The plumbing construction is to take two disk bundles $W(n)$ and $W(m)$, and trivialize neighborhoods of points p and p' in the bases. We then quotient $W(n) \sqcup W(m)$ by identifying the two copies of $D^2 \times D^2$ in a way which switches the base and fiber directions. See Figure 4.1.

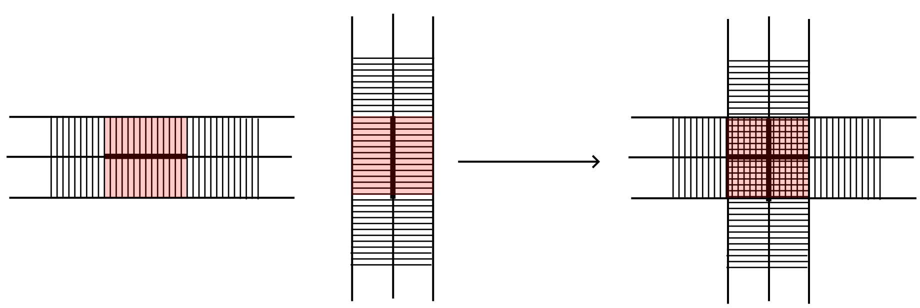

Figure 4.1. Plumbing

Given the tree G, we take the disk bundles and then perform the plumbing construction for each edge of G. This yields a rather complicated 3 and 4-manifold.

Remark 4.2. There is an alternate description in terms of Dehn surgery. For each vertex we take an unknot. For each edge of G we add a clasp. Call the link L_G. We give each unlink the framing corresponding to its weight. Then $S^3(L_G) \cong Y(G)$.

4.2. **Characteristic vectors and intersection forms.**

A *characteristic vector* of a 4-manifold W is an element $K \in H^2(W)$ such that if $s \in H_2(W)$, then

$$K(s) + s \cdot s \equiv 0 \pmod 2.$$

Example 4.3. If $V = (\pm 1)$, then the characteristic vectors are the odd integers. If V is a direct sum of copies of (± 1), then the characteristic vectors are tuples of odd integers.

In Heegaard Floer theory, there is an important algebraic topological construction called a Spin^c structure. Given a 3-manifold or 4-manifold X, there is a set $\mathrm{Spin}^c(X)$. This set is an affine space over $H^2(X;\mathbb{Z})$, and in fact, is non-canonically isomorphic to $H^2(X;\mathbb{Z})$ as an affine space. Note $\mathrm{Spin}^c(X)$ does not naturally have an additive structure (i.e. group structure) but rather has an action of $H^2(X;\mathbb{Z})$. Additionally, there is a chern class map

$$c_1\colon \mathrm{Spin}^c(X) \to H^2(X;\mathbb{Z}).$$

This map is *not* an isomorphism, but it does respect the affine structure via the equation

$$c_1(\mathfrak{s}+h) = c_1(\mathfrak{s}) + 2h. \tag{4.1}$$

Lemma 4.4. *Characteristic vectors of a 4-manifold W and in bijection with Chern classes of* Spin^c *structures on W.*

Exercise 4.5. Prove the above using the fact that $c_1(\mathfrak{s})$ has mod 2 reduction equal to $w_2(TW)$ and the fact that every element in $H_2(W;\mathbb{Z})$ can be represented by a smoothly embedded surface.

Remark 4.6. If G is a plumbing tree, we can form a matrix $\Lambda_G = (\lambda_{v,v'})_{v,v'\in V(G)}$ as follows. We set $\lambda_{v,v} = w(v)$, and we set $\lambda_{v,v'} = 1$ if v and v' are connected by an edge. Set $\lambda_{v,v'} = 0$ otherwise. It is straightforward to prove that $H_1(Y(G))$ has a presentation as $\mathbb{Z}^n/\operatorname{im}\Lambda_G$. Also, one can show that if Λ_G is non-singular, then $|\det(\Lambda_G)|$ is the number of elements in $H_1(Y(G))$. This is the same as the number of Spin^c structures. (If $\det(\Lambda_G) = 0$, then there are infinitely many elements).

Remark 4.7. For a plumbing tree G, we have $\mathrm{Spin}^c(Y(G)) \cong \mathrm{Char}(X(G))/H_2(X(G))$.

Note that $H_2(X(G)) \cong \mathbb{Z}^n$, where n is the number of vertices. By the universal coefficient theorem, we can identify $H^2(X(G)) \cong \mathrm{Hom}(H_2(X(G)),\mathbb{Z})$ (since $H_1(X(G)) \cong 0$). Furthermore, there is a natural inclusion $H_2(X(G)) \to \mathrm{Hom}(H_2(X(G)),\mathbb{Z}) = H^2(X(G))$.

4.3. **Motivation for the lattice complex.** Given a plumbed 3-manifoly $Y(G)$, we have observed that there is a canonical 4-manifold $W(G)$ which bounds $Y(G)$. We assume that $W(G)$ is negative definite.

We can imagine "probing" $HF^-(Y(G))$ using the cobordism maps for $W(G)$. For each $\mathfrak{s} \in \mathrm{Spin}^c(W(G))$, we get a map

$$F_{W,\mathfrak{s}}\colon HF^-(S^3) \to HF^-(Y(G)).$$

Since $HF^-(S^3) \cong \mathbb{F}[U]$, and $F_{W,\mathfrak{s}}$ commutes with U, this is the same as an element of $HF^-(Y(G))$. Since $H^2(W(G))$ is torsion free, the map

$$c_1\colon \mathrm{Spin}^c(W(G)) \to \mathrm{Char}(W(G))$$

is a bijection, so we find it cleaner work with characteristic vectors instead of Spin^c structures.

Therefore, we have an element $[K]$ in $HF^-(Y(G))$ for each $K \in \mathrm{Char}(W(G))$, which is the image of 1 under the corresponding cobordism map. There are two questions to ask:

(1) Do the elements $[K]$ span all of $HF^-(Y(G))$?
(2) What are the relations between $[K]$ and $[K']$ for different characteristic vectors?

We can partially answer the second one, as follows. Note that $W(G)$ has a large number of embedded spheres contained in it, i.e. one for each vertex of G. The boundary of a neighborhood of each spheres is a lens space $L(w,1)$, where w is the weight of the sphere (Euler number). Recall that $L(w,1)$ is an L-space, i.e. has homology

$$\bigoplus_{|w|}\mathbb{F}[U],$$

where each $\mathbb{F}[U]$ is concentrated in a different Spin^c structure.

Let E be one of these spheres, and let $N=N(E)$ be a tubular neighborhood. Note N is a disk bundle over sphere with Euler number $E\cdot E$.

Note that we can factor the cobordism map

$$F_{W,\mathfrak{s}}=F_{W\setminus N,\mathfrak{s}|_{W\setminus N}}\circ F_{N,\mathfrak{s}|_N}.$$

We consider the effect of changing $\mathfrak{s}$ to $\mathfrak{s}\pm PD[E]$. Note that this has the effect of adding $\pm 2E$ to K by Equation (4.1).

Note that $F_{W\setminus N,\mathfrak{s}|_{W\setminus N}}$ is unchaged by adding $PD[E]$ to $\mathfrak{s}$. On the other hand $F_{N,\mathfrak{s}|_N}$ may be chanced, though the restriction of $\mathfrak{s}$ to ∂N is unchanged. In particular

$$F_{N,\mathfrak{s}}\doteq F_{N,\mathfrak{s}-PD[E]}$$

where $\doteq$ means equality up to multiplication by some power of U.

Ergo, $[K]$ and $[K+2E]$ must be equal after multiplying by some large powers of U. Typically we write just $v\in H_2$ instead of E_v, to abbreviate.

Note too that the gradings of the elements $[K]$ and $[K']$ are determined by the formula

$$\mathrm{gr}(K)=\frac{K^2-3\sigma(K)-2\chi(K)}{4}$$

so the powers of U are determined from homological considerations. Here, we are implicitly factoring K, which lies in $H^2(X(G))$, into $H^2(X(G),Y(G))$ using the fact that $Y(G)$ is a $\mathbb{Q}HS^3$. Then K^2 denotes the cup product evaluated on the fundamental class $[X(G),Y(G)]$.

We can formalize this a bit as follows. If $K\in\mathrm{Char}(W)$ and $v\in V(G)$, we let

$$\chi_K(v)=\frac{K(v)+v\cdot v}{2}.$$

Remark 4.8. Note that

$$\chi_K(v)=\frac{\mathrm{gr}(K+2v)-\mathrm{gr}(K)}{2}.$$

If $\chi_K(v)\geq 0$, then we have the relation

$$U^{\chi_K(v)}\cdot[K+2v]=[K]$$

and if $\chi_K(v)\leq 0$, then we have the relation

$$[K+2v]=U^{-\chi_K(v)}[K].$$

Therefore we may make a guess as to $HF^-(Y(G))$ by encoding just the above generators and relations. Namely, we define $\mathbb{H}^-(G)$ to be generated over $\mathbb{F}[U]$ by all K, subject only to the relations above. I.e.

$$\mathbb{H}^-(G):=\frac{\bigoplus_{K\in\mathrm{Char}(W)}\mathbb{F}[U]}{R}$$

where R denotes the above relations.

Theorem 4.9 ([OS03] [Ném05]). *If G is negative definite and "suitably nice" then $\mathbb{H}^-(G) \cong HF^-(Y(G))$.*

The "suitably nice" condition concerns the combinatorics of the graph. The condition that Ozsváth and Szabó used was the following: we say a vertex v is *bad* if

$$m(v) > -\operatorname{val}(v).$$

(Here m is the multiplicity assumed to be negative). They proved the above theorem when there is at most one bad vertex.

Némethi interpretted (and generalized) this condition, showing that the bad vertex condition has a natural interpretation in terms of normal surface singularities and algebraic geometry.

As philosophical motivation, we can view the number of bad vertices as giving a measure of complexity of the graph. Write $\mathcal{G}_{\le n}$ for the set of negative definite trees with at most n bad vertices. Then we have inclusions

$$\mathcal{G}_{\le 0} \subseteq \mathcal{G}_{\le 1} \subseteq \mathcal{G}_{\le 2} \subseteq \cdots .$$

Having more bad vertices in fact means that the 3-manifold $Y(G)$ and $HF^-(Y(G))$ get more complicated.

For example, it turns out that if $G \in \mathcal{G}_{\le 0}$ then $Y(G)$ is an L-space. If $G \in \mathcal{G}_{\le 1}$, the $HF^-(Y(G), \mathfrak{s})$ is supported in a single mod 2 grading for each $\mathfrak{s}$.

Remark 4.10. [Ném17] Némethi proves a modification of the above which is an if and only if statement. Némethi considers a different condition called *rationality*, from algebraic geometry. He shows that if there are no bad vertices, then G is rational. Furthermore, he shows that G is rational if and only if $Y(G)$ is an L-space.

4.4. **Computations.** In general, the above theorem reduces the computation of $HF^-(Y(G))$ to a combinatorial problem, though considerable additional work is needed to produce an actual computation. See [Ném05] [Ném08] for further computational techniques. For these rather introductory notes, we will content ourselves with understanding the case where the plumbing graph has a single vertex v, with weight $n < 0$.

We have $H_2(X(G)) = \mathbb{Z}$ and $H^2(X(G)) = \operatorname{Hom}(\mathbb{Z}, \mathbb{Z}) = \mathbb{Z}$. If $j \in \mathbb{Z}$, write $[j]$ for the corresponding element in $H^2(X(G))$, normalized so that $K_{[j]}(v) = j$. We have $K_{[j]}$ is characteristic iff

$$K_{[j]}(v) + v^2 \equiv 0, \quad \text{ie} \quad j + n \equiv 0 \pmod 2.$$

So K is characteristic iff the corresponding $j \in \mathbb{Z}$ has the same parity as n.

For the purposes of demonstration, consider the case that $n = -2$. Then characteristic vectors are identified with even integers $j \in 2\mathbb{Z}$. We consider the adjunction relations. We compute that

$$\chi_{[j]}(v) = \frac{j+n}{2}.$$

The characteristic vectors are partitioned into two classes in $\operatorname{Char}(X(G))/H_2(X(G)) \cong \operatorname{Spin}^c(Y(G))$, namely the class of $[0]$ and the class of $[2]$.

Note that $[j] + 2v = [j-4]$.

The first few adjunction relations read

$$U[-4] = [0] \quad \text{and} \quad U^3[-8] = [-4].$$

Filling in the rest, we get

$$\cdots = UU^3[-8] = U[-4] = [0] = U[4] = UU^3[8] = UU^3U^5[12] = \cdots$$

The other Spin^c structure is

$$\cdots = U^2U^4[-10] = U^2[-6] = [-2] = [2] = U^2[6] = U^2U^4[10] = \cdots$$

In particular, HF^- is generated by [0] and [2] and is thus $\mathbb{F}[U] \oplus \mathbb{F}[U]$.

References

[Aur13] Denis Auroux, *A beginner's introduction to fukaya categories*, 2013. e-print, arXiv 1301.7056.

[BLZ22] Maciej Borodzik, Beibei Liu, and Ian Zemke, *Lattice homology, formality, and plumbed L-space links*, 2022. e-print, arXiv 2210.15792 [math.GT].

[GN16] Eugene Gorsky and András Némethi, *Links of plane curve singularities are L-space links*, Algebr. Geom. Topol. **16** (2016), no. 4, 1905–1912.

[Hed09] Matthew Hedden, *On knot Floer homology and cabling. II*, Int. Math. Res. Not. IMRN **12** (2009), 2248–2274.

[Mos71] Louise Moser, *Elementary surgery along a torus knot*, Pacific J. Math. **38** (1971), 737–745.

[Ném05] András Némethi, *On the Ozsváth-Szabó invariant of negative definite plumbed 3-manifolds*, Geom. Topol. **9** (2005), 991–1042.

[Ném08] ———, *Lattice cohomology of normal surface singularities*, Publ. Res. Inst. Math. Sci. **44** (2008), no. 2, 507–543.

[Ném17] ———, *Links of rational singularities, L-spaces and LO fundamental groups*, Invent. Math. **210** (2017), no. 1, 69–83.

[OS03] Peter Ozsváth and Zoltán Szabó, *On the Floer homology of plumbed three-manifolds*, Geom. Topol. **7** (2003), 185–224.

[OS05] Peter Ozsváth and Zoltán Szabó, *On knot Floer homology and lens space surgeries*, Topology **44** (2005), no. 6, 1281–1300.

[OS08] Peter S. Ozsváth and Zoltán Szabó, *Knot Floer homology and integer surgeries*, Algebr. Geom. Topol. **8** (2008), no. 1, 101–153.

[Sei08] Paul Seidel, *Fukaya categories and Picard-Lefschetz theory*, Zurich Lectures in Advanced Mathematics, European Mathematical Society (EMS), Zürich, 2008.

[Wei94] Charles A. Weibel, *An introduction to homological algebra*, Cambridge Studies in Advanced Mathematics, vol. 38, Cambridge University Press, Cambridge, 1994.