

BRAID LOOPS WITH INFINITE MONODROMY ON THE LEGENDRIAN CONTACT DGA

ROGER CASALS AND LENHARD NG

ABSTRACT. We present the first examples of elements in the fundamental group of the space of Legendrian links in $(\mathbb{S}^3, \xi_{\text{st}})$ whose action on the Legendrian contact DGA is of infinite order. This allows us to construct the first families of Legendrian links that can be shown to admit infinitely many Lagrangian fillings by Floer-theoretic techniques. These new families include the first known Legendrian links with infinitely many fillings that are not rainbow closures of positive braids, and the smallest Legendrian link with infinitely many fillings known to date. We discuss how to use our examples to construct other links with infinitely many fillings, and in particular give the first Floer-theoretic proof that Legendrian (n, m) torus links have infinitely many Lagrangian fillings if $n \geq 3, m \geq 6$ or $(n, m) = (4, 4), (4, 5)$. In addition, for any given higher genus, we construct a Weinstein 4-manifold homotopic to the 2-sphere whose wrapped Fukaya category can distinguish infinitely many exact closed Lagrangian surfaces of that genus in the same smooth isotopy class, but distinct Hamiltonian isotopy classes. A key technical ingredient behind our results is a new combinatorial formula for decomposable cobordism maps between Legendrian contact DGAs with integer (group ring) coefficients.

Contents

1.	Introduction	2
	1.1. Context	3
	1.2. Main Results	5
2.	Legendrian Links and ϑ-loops	9
	2.1. Legendrian links, exact Lagrangian cobordisms, and fillings	9
	2.2. Legendrian links associated to positive braids	11
	2.3. A class of Legendrian (-1) -closures	16
	2.4. The purple-box Legendrian loop	18
3.	Legendrian Contact DGAs and Cobordism Maps	20
	3.1. The Legendrian contact DGA	20
	3.2. Link automorphisms	23
	3.3. The geometric map induced by an exact Lagrangian cobordism	24
	3.4. Signs and functoriality of the cobordism map	25
	3.5. System of augmentations for a decomposable filling	27
	3.6. The systems of augmentations agree	30
4.	Cobordism Maps for Elementary Cobordisms	34
	4.1. The cobordism map for a Legendrian isotopy	34

2010 *Mathematics Subject Classification*. Primary: 53D10. Secondary: 53D15, 57R17.

4.1.1. Base point moves	35
4.1.2. Reidemeister III moves	35
4.1.3. Reidemeister II moves	36
4.2. The cobordism map for a saddle cobordism	37
4.3. Assembling elementary cobordism maps	40
5. Legendrian Contact DGA and Cobordism Maps for (-1)-closures	42
5.1. The DGA of the (-1) -closure of an admissible braid	42
5.2. ϑ -monodromy action on the DGA	43
5.3. The Kálmán loop	46
5.4. The saddle cobordism map for (-1) -closures	47
6. Proof of Infinitely Many Fillings	49
6.1. The argument	49
6.2. Augmentations for $\Lambda(\tilde{D}_4)$	49
6.3. Three Variations on the Affine D_4 -braid	53
6.4. Monodromy for the Braids Λ_1 and Λ_2	55
6.5. The general case: $\Lambda(\tilde{D}_n)$ and Λ_n	56
7. Proof of Corollaries and Concluding Remarks	58
7.1. Smooth isotopy class of Lagrangian fillings	58
7.2. Aug-infinite Legendrian links and cobordisms	60
7.3. Lagrangian surfaces in Weinstein 4-manifolds	62
Appendix A. The Cobordism Map for an Elementary Saddle Cobordism	67
A.1. Formula for the saddle cobordism map	67
A.2. Splashes and diagonal automorphisms	68
A.3. Proof of Proposition A.1	73
A.4. Proof of Proposition 4.8	77
References	80

1. INTRODUCTION

In this article, we construct Legendrian loops for several families of Legendrian links in the standard contact 3-sphere $(\mathbb{S}^3, \xi_{\text{st}})$ and show that their monodromy action on their Legendrian contact DGA is of infinite order. These are the first examples of such a Floer-theoretic infinite order, in sharp contrast with the known finite order DGA action of all previously studied loops. We provide several new consequences of these results, including the first known examples of Legendrian links with infinitely many Lagrangian fillings which are *not*

the rainbow closure of a positive braid¹, and can be distinguished via Floer theory.² These Lagrangian fillings are all smoothly isotopic, but their Hamiltonian isotopy classes are all distinct. One of these new Legendrian links has 2 components and, with its Lagrangian fillings being of genus 1, is arguably the smallest known Legendrian link to date, in terms of genus and components, with infinitely many Lagrangian fillings. In addition, for any given genus $g \geq 2$, we construct Weinstein 4-manifolds homotopic to the 2-sphere whose wrapped Fukaya categories can distinguish infinitely many (Hamiltonian isotopy classes of) exact closed Lagrangian surfaces of that genus, all in the same smooth type. Finally, we show how to Floer-theoretically detect the existence of infinitely many Lagrangian fillings for the Legendrian (n, m) torus links of maximal Thurston–Bennequin number (“max-tb”), with $n \geq 3, m \geq 6$ and $(n, m) = (4, 4), (4, 5)$, and many other Legendrian links, by using the Legendrian DGA.³

The manuscript also develops technical results on the Legendrian contact DGA, of independent interest, needed for our argument. In particular, we present a combinatorial model for computing DGA morphisms associated to decomposable Lagrangian cobordisms L , where the morphisms are enhanced over *integer* group ring coefficients. We show that this is isomorphic to the abstract enhancement previously developed by Karlsson, thus proving invariance and allowing us to perform explicit computations over $\mathbb{Z}[H_1(L)]$. This integrally enhanced package is then used to prove the above Floer-theoretical results concerning infinitely many Lagrangian fillings.

1.1. Context. Legendrian links in contact 3-manifolds [Ben83, Ad90] are instrumental in the study of 3-dimensional contact geometry [OS04, Gei08]. The study of their Lagrangian fillings yields non-trivial DGA representations of the Legendrian contact DGA associated to any Legendrian link, which themselves are effective invariants for distinguishing Legendrian representatives in the same smooth type [Che02, Ng03, Siv11]. In particular, Floer theory has provided far-reaching methods to address questions on Legendrian links; for instance, along the lines of this paper, see [EP96, Etn03, Kál05, Cha10].

Recently, the first examples of Legendrian links in $(\mathbb{S}^3, \xi_{\text{st}})$ which admit *infinitely* many Lagrangian fillings in $(\mathbb{D}^4, \lambda_{\text{st}})$ were discovered [CG21]. Indeed, [CG21, Corollary 1.5] shows that the max-tb Legendrian (n, m) -torus link $\Lambda(n, m)$ admits infinitely many Lagrangian fillings if $n \geq 3, m \geq 6$ or $(n, m) = (4, 4), (4, 5)$. The method of proof itself relies on the theory of microlocal sheaves, and it remained unclear whether the existence of infinitely many Lagrangian fillings, even for one Legendrian link, could also be proven via Floer-theoretic methods. It also remained unknown whether (typically smaller) links which were not rainbow closures of positive braids – from which the current sheaf methods do not apply – could actually admit infinitely many Lagrangian fillings.

(i) First, we show that the Legendrian DGA detects infinitely many fillings and it does so for new Legendrian links (including links that are not the rainbow closure of a positive braid). In fact, we significantly improve on [CG21, Corollary 1.5] by showing that simpler classes of Legendrian braids already admit infinitely many exact Lagrangian fillings, and doing so Floer-theoretically. For instance, the family of Legendrian braids of affine D_n -type depicted in Figure 1 (right) is one such class. This also gives an alternative Floer-theoretical proof that the torus links in [CG21, Corollary 1.5] admit infinitely many Lagrangian fillings.

¹Previously known methods to build infinite Lagrangian fillings, including the techniques from microlocal sheaf theory, do not apply in this general setting.

²Note that, before this manuscript, none of the infinite Lagrangian fillings in [CG21] or [CZ21] was known to be distinguished via Floer theory.

³In all cases being considered, the max-tb condition is a necessary condition on the Legendrian links in order to admit an embedded exact Lagrangian filling, see e.g. [Cha10].




FIGURE 1. The family of Legendrian links $\Lambda_n \subset (\mathbb{S}^3, \xi_{\text{st}})$, $n \geq 1$, on the left. The Legendrian links of affine D_n -type are depicted on the right, $n \geq 4$. All of these have infinitely many fillings. The boxes indicate a series of positive crossings.

Note that, since the appearance of [CG21], the articles [CZ21, GSW20a, GSW20b] have also continued to develop various cluster and sheaf-theoretic methods that detect infinitely many Lagrangian fillings for a Legendrian link $\Lambda \subset (\mathbb{S}^3, \xi_{\text{st}})$. Nevertheless, all these techniques are currently only effective at studying Legendrian links which are positive braids, i.e. when $\Lambda \subset (\mathbb{S}^3, \xi_{\text{st}})$ admits a Legendrian front given by the rainbow closure of a positive braid, and do *not* apply to several of our smallest links. In contrast, the Floer-theoretic argument we develop *also* applies to certain Legendrian links $\Lambda \subset (\mathbb{S}^3, \xi_{\text{st}})$ which are *not* the rainbow closure of positive braids. For instance, we show that each of the Legendrian links Λ_n , $n \in \mathbb{N}$, depicted in Figure 1 (left) admits infinitely many Lagrangian fillings. For $n = 1$, this yields a Legendrian link Λ_1 which is not the rainbow closure of a positive braid because it contains a $tb = -3$ stabilized unknot component.

(ii) Second, the existence of infinitely many Lagrangian fillings for our families of Legendrian links $\Lambda \subset (\mathbb{S}^3, \xi_{\text{st}})$ is deduced from a stronger result on Legendrian loops, Theorem 1.1, as we explain shortly. In particular, we provide the first examples of Legendrian loops whose induced monodromy action on the Legendrian contact DGA has infinite order. In addition, we present the first example of a Weinstein 4-manifold homotopic to the 2-sphere with infinitely many Hamiltonian isotopy classes of exact Lagrangian surfaces of genus 2 (and no Lagrangian 2-spheres nor exact Lagrangian tori). This is part of the family of Weinstein 4-manifolds in Corollaries 1.6 and 1.7, which construct such Weinstein 4-manifolds for all genera $g \geq 2$. Note that, at the level of smooth topology, the concatenation of these Legendrian loops with any decomposable Lagrangian filling does *not* change the smooth type of the Lagrangian filling. Thus, we can use these Legendrian loops to produce infinitely many Lagrangian fillings (and surfaces in Weinstein 4-manifolds) which are distinct up to Hamiltonian isotopy, but these surfaces are all *smoothly* isotopic.

(iii) Third, at a technical level, we study the lifts of the DGA maps induced by exact Lagrangian cobordisms to \mathbb{Z} -coefficients, which is required to argue the infinite order in our argument. This is interesting on its own, as it provides correct signs for Floer theoretical invariants, such as augmentations, and it is a necessary ingredient for the study of cluster structures⁴ on augmentation varieties and their holomorphic symplectic structures, as this requires Floer theory in characteristic 0. In particular, these results from this manuscript are used in the recent article [CGGS20] to construct a holomorphic symplectic structure on the augmentation varieties associated to Legendrian positive braids.

⁴In characteristic different from 2. In particular, the correct signs are needed for arguing in characteristic 0, the most studied case in cluster theory.




FIGURE 2. On the left, Lagrangian projection of the Legendrian links Λ_n , $n \geq 1$, and the purple box, which contains n positive crossings. The purple-box Legendrian loop $\vartheta : \mathbb{S}^1 \rightarrow \mathcal{L}(\Lambda_n)$ is depicted by the dashed purple trajectory. On the right, Lagrangian projection of the Legendrian link $\Lambda(\tilde{D}_n)$, $n \geq 4$, and the purple box, which contains $(n-2)$ positive crossings. The purple-box Legendrian loop $\vartheta : \mathbb{S}^1 \rightarrow \mathcal{L}(\Lambda(\tilde{D}_n))$ is also illustrated by the dashed purple trajectory.

1.2. Main Results. Let β be a positive braid, representing an element in the N -stranded positive braid monoid Br_N^+ , $N \in \mathbb{N}$. We can associate a Legendrian link $\Lambda(\beta) \subset (\mathbb{S}^3, \xi_{\text{st}})$ to β such that $\Lambda(\beta)$ is topologically the (-1) -framed closure of β : this is achieved by placing β in a standard contact neighborhood of the standard Legendrian unknot in \mathbb{S}^3 of Thurston–Bennequin number $tb = -1$. See Figure 3 for a depiction of $\Lambda(\beta)$, where throughout this paper we will describe Legendrian links through their front and/or Lagrangian projections (see Section 2 for a review). We now define the Legendrian links that we will study in this paper.




FIGURE 3. The general front for the Legendrian links $\Lambda(\beta)$, on the left, and the specific example of the Legendrian link $\Lambda(\beta_{11})$, on the right.

By definition, the \tilde{D}_n -Legendrian link is the Legendrian $\Lambda(\beta(\tilde{D}_n)) \subset (\mathbb{S}^3, \xi_{\text{st}})$ associated to

$$\beta(\tilde{D}_n) = (\sigma_2 \sigma_1 \sigma_3 \sigma_2 \sigma_2 \sigma_3 \sigma_1 \sigma_2) \sigma_1^{n-4} \Delta^2, \quad n \geq 4,$$

where $\Delta = \sigma_1(\sigma_2 \sigma_1)(\sigma_3 \sigma_2 \sigma_1)$ is the 4-stranded half-twist. Figure 1 (right) shows a Legendrian front projection for $\Lambda(\tilde{D}_n)$, and the terminology will be explained in Section 2.⁵

Similarly, the Legendrian link $\Lambda_n \subset (\mathbb{S}^3, \xi_{\text{st}})$ is the Legendrian link $\Lambda_n = \Lambda(\beta_n)$ associated to the braid word

$$\beta_n = (\sigma_2 \sigma_1 \sigma_1 \sigma_2)^3 \sigma_1^n, \quad n \geq 1.$$

Figure 1 (left) shows a Legendrian front projection for Λ_n . These are two distinct families of Legendrian links, with the exception of the accidental Legendrian isotopy $\Lambda_2 \cong \Lambda(\tilde{D}_5)$.

⁵The \tilde{D}_n -Legendrian should be read as the *affine* D_n -Legendrian.

Finally, we will also consider the Legendrian links associated to the following braids:

$$\beta_{ab} = (\sigma_2\sigma_1\sigma_3\sigma_2)^4\sigma_3^a\sigma_1^b, \quad a, b \in \{1, 2\}.$$

Note that $\Lambda(\beta_{22})$ is Legendrian isotopic to $\Lambda(\tilde{D}_4)$. See Figure 3 (right) for a drawing of $\Lambda(\beta_{11})$. Following the above Dynkin-diagram notation, $\Lambda(\beta_{11})$ can also be referred to as the $\Lambda(\tilde{A}_{2,1})$ -Legendrian link. From now onwards, we denote by

$$\mathcal{H} = \{\Lambda_n\}_{n \geq 1} \cup \{\Lambda(\tilde{D}_m)\}_{m \geq 4} \cup \{\Lambda(\beta_{11}), \Lambda(\beta_{12}), \Lambda(\beta_{21})\}$$

the set-theoretic union of the Legendrian links in the Λ_n and $\Lambda(\tilde{D}_n)$ families described above and the three Legendrians links $\Lambda(\beta_{11}), \Lambda(\beta_{12}), \Lambda(\beta_{21})$. The Legendrian links in \mathcal{H} allow us to tackle a wide range of additional Legendrian links, thanks to Corollary 1.3 below. This includes torus links, as in Corollary 1.4, and the knots discussed in Section 7, see Remark 1.5 below.

Let $\mathcal{L}(\Lambda)$ be the space of Legendrian links isotopic to the Legendrian link $\Lambda \subset (\mathbb{S}^3, \xi_{\text{st}})$, with base point an arbitrary but fixed Legendrian representative. In Section 2, for each of the links $\Lambda \in \mathcal{H}$, we will define a certain loop ϑ of Legendrians based at Λ : that is, a continuous map $\vartheta : (\mathbb{S}^1, \text{pt}) \rightarrow (\mathcal{L}(\Lambda), \Lambda)$. For instance, for the Legendrians in Figure 2, the loop arises from moving the purple box around the link in the manner depicted. We will refer to this Legendrian loop ϑ as the purple-box Legendrian loop.

The graph of the Legendrian loop ϑ produces an exact Lagrangian concordance L_ϑ in the symplectization of $(\mathbb{S}^3, \xi_{\text{st}})$ from Λ to itself. Given any filling $L \subset (\mathbb{D}^4, \lambda_{\text{st}})$ of Λ , which we can view as an exact Lagrangian cobordism from the empty link to Λ , we can concatenate L with any number of copies of L_ϑ to produce an infinite family of fillings

$$L \# L_\vartheta^n, \quad n \in \mathbb{N},$$

of Λ . What we will show is that for $\Lambda \in \mathcal{H}$, we can choose a filling L of Λ such that all of these fillings $L \# L_\vartheta^n$ are distinct.

As discussed earlier, our method of proof involves the Legendrian contact DGA \mathcal{A}_Λ of Λ , which is an invariant of the Legendrian isotopy class of $\Lambda \subset (\mathbb{S}^3, \xi_{\text{st}})$, up to stable tame DGA isomorphism. The concordance L_ϑ induces a DGA isomorphism

$$\mathcal{A}(L_\vartheta) : \mathcal{A}_\Lambda \rightarrow \mathcal{A}_\Lambda$$

while the filling L induces a DGA morphism (“augmentation”)

$$\varepsilon_L : \mathcal{A}_\Lambda \rightarrow (\mathbb{Z}[H_1(L)], 0),$$

where $(\mathbb{Z}[H_1(L)], 0)$ is the DGA with trivial differential, concentrated in degree 0. Functoriality then implies that the filling $L \# L_\vartheta^n$ induces the augmentation $\varepsilon_L \circ \mathcal{A}(L_\vartheta)^n$. To distinguish the fillings $L \# L_\vartheta^n$ from each other, we will distinguish the augmentations $\varepsilon_L \circ \mathcal{A}(L_\vartheta)^n$, even allowing for different choices of local systems on the fillings.

To be precise, we say that the ϑ -orbit of the augmentation ε_L is *entire* if for any $k, l \in \mathbb{N}$ distinct, there is no automorphism $\varphi \in \text{Aut}(\mathbb{Z}[H_1(L)])$ such that

$$\varphi(\varepsilon_L \circ \mathcal{A}(L_\vartheta)^k) = \varepsilon_L \circ \mathcal{A}(L_\vartheta)^l : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}[H_1(L)].$$

The first result in our article is the following:

Theorem 1.1. *Let $\Lambda \in \mathcal{H}$ be a Legendrian link. The purple-box Legendrian loop $\vartheta : \mathbb{S}^1 \rightarrow \mathcal{L}(\Lambda)$ induces a DGA map $\mathcal{A}(L_\vartheta) : \mathcal{A}(\Lambda) \rightarrow \mathcal{A}(\Lambda)$ of infinite order. In fact, there exists an exact Lagrangian filling $L \subset (\mathbb{D}^4, \lambda_{\text{st}})$ such that the ϑ -orbit of the corresponding augmentation $\varepsilon_L : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}[H_1(L)]$ is entire.*

To our knowledge, Theorem 1.1 presents the first Legendrian loops which induce an *infinite* order action on the augmentations of a Legendrian contact DGA \mathcal{A}_Λ . Our work is a spiritual

successor to the work of T. Kálmán [Kál05], who studied Legendrian loops for positive torus links $\Lambda(n, m)$ whose induced action on $\mathcal{A}(\Lambda(n, m))$ has finite order $(n + m)$.

Theorem 1.1 implies the following:

Corollary 1.2. *Let $\Lambda \in \mathcal{H}$. Then the purple-box Legendrian loop ϑ generates an infinite subgroup $\mathbb{Z}\langle\vartheta\rangle \subset \pi_1(\mathcal{L}(\Lambda))$. In addition, the graph of the Legendrian loop ϑ produces a Lagrangian self-concordance of Λ which has infinite order as an element of the Lagrangian concordance monoid based at Λ .*

Let us now focus on Lagrangian fillings. Theorem 1.1 implies that each of the Legendrian links $\Lambda \in \mathcal{H}$ admits infinitely many Lagrangian fillings, up to Hamiltonian isotopy. More precisely, there exists a countably infinite collection $\{L_i\}_{i \in \mathbb{N}}$ of oriented embedded exact Lagrangian fillings $L_i \subset (\mathbb{D}^4, \lambda_{\text{st}})$ of the Legendrian link Λ in the boundary $\mathbb{S}^3 = \partial\mathbb{D}^4$ such that all L_i are smoothly isotopic for $i \in \mathbb{N}$, relative to a neighborhood of the boundary Λ , but none of the L_i are Hamiltonian isotopic to each other; that is, if $i \neq j$, there exists no compactly supported Hamiltonian isotopy $\{\varphi_t\} \in \text{Ham}^c(\mathbb{D}^4, \lambda_{\text{st}})$, $\varphi_0 = \text{Id}$, such that $\varphi_1(L_i) = L_j$.

We note that among the Legendrian links in \mathcal{H} , four links— Λ_1 , $\Lambda(\beta_{11})$, $\Lambda(\beta_{12})$, and $\Lambda(\beta_{21})$ —have a component which is a stabilized unknot with Thurston–Bennequin number -3 . (In fact $\Lambda(\beta_{11})$ has two such components.) It follows that none of these four links is the rainbow closure of a positive braid. We emphasize that the methods developed in [CG21, CZ21, GSW20a, GSW20b] for the detecting of infinitely many Lagrangian fillings only apply to rainbow closures of positive braids, and thus our Floer-theoretic techniques provide new results that we currently do not know how to address through cluster algebras [GSW20a, GSW20b] or the study of microlocal sheaves [CG21, CZ21].⁶

We can use Legendrian links with infinitely many fillings to produce other Legendrian links with infinitely many fillings. Roughly speaking, if there is an exact Lagrangian cobordism from Λ_- to Λ_+ and Λ_- has infinitely many fillings, then Λ_+ does as well. (We only prove this statement subject to some important hypotheses; see Proposition 7.5 for the precise result.) In particular, we have the following consequence of Theorem 1.1.

Corollary 1.3 (see Proposition 7.5). *Let $\Lambda_0, \Lambda \subset (\mathbb{S}^3, \xi_{\text{st}})$ be Legendrian links with Λ_0 in the list \mathcal{H} , and suppose that there is a Lagrangian cobordism from Λ_0 to Λ consisting of a sequence of saddle moves at contractible Reeb chords of degree 0. Then the Legendrian link Λ admits infinitely many exact Lagrangian fillings, distinct up to Hamiltonian isotopy.*

As a special case, since there are such cobordisms to the max-tb Legendrian (n, m) torus links $\Lambda(n, m)$ from Λ_1 for $n = 3, m \geq 6$, and from $\Lambda(\tilde{D}_4)$ for $n, m \geq 4$, we recover the following result of [CG21].

Corollary 1.4 ([CG21]). *The Legendrian torus links $\Lambda(n, m)$ each admit infinitely many exact Lagrangian fillings if $n \geq 3, m \geq 6$ or $(n, m) = (4, 4), (4, 5)$.*

Remark 1.5. As we will discuss in Section 7.2, among the universe of Legendrian links with infinitely many fillings, a sensible notion of “simplicity” is given by the Thurston–Bennequin number, or equivalently the sum $2g + m$, where g is the genus of an exact Lagrangian filling and m is the number of connected components of the link: the smaller $2g + m$ is, the simpler the link is. Among the Legendrian links that we can prove have infinitely many fillings, the simplest by this measure is $\Lambda(\beta_{11})$, which has $(m, g) = (2, 1)$ and thus $2g + m = 4$.

If we focus on Legendrian *knots*, rather than Legendrian *links*, Corollary 1.3 implies that, for instance, the knot types 10_{139} , $m(10_{145})$, $m(10_{152})$, 10_{154} , and $m(10_{161})$ all have Legendrian representatives with infinitely many fillings; see Proposition 7.7. Among these, the simplest

⁶By Corollary 1.3, we can in fact construct an infinite family of links with infinitely many fillings that are not the rainbow closure of a positive braid: $\Lambda((\sigma_2\sigma_1\sigma_3\sigma_2)^4\sigma_3^a\sigma_1)$ for $n \geq \mathbb{N}$.

is $m(10_{145})$, with $g = 2$ and $2g + m = 5$. Two of these knots, 10_{139} and $m(10_{152})$, are positive braid closures and indeed their Legendrian representatives are rainbow closures of positive braids. We remark that the only other knots with crossing number ≤ 10 that are positive braid closures are the torus knots $T(2, 3)$, $T(2, 5)$, $T(2, 7)$, $T(3, 4)$, $T(2, 9)$, and $T(3, 5)$; it is conjectured that the (max-tb) Legendrian representatives of each of these knots has finitely many fillings [Cas21, Conjecture 5.1]. \square

The above results on Lagrangian fillings also have consequences in the study of Stein surfaces. For each $g \in \mathbb{N}$ and $g \geq 6$, the article [CG21] gave the first examples of Stein surfaces homotopic to the 2-sphere \mathbb{S}^2 with infinitely many Hamiltonian isotopy classes of embedded exact Lagrangian surfaces of genus g (and none of genus less than g). The lower bound was recently improved to $g \geq 4$ in [GSW20b]. In the present work, we can further improve this bound:

Corollary 1.6. *Let $g \in \mathbb{N}$ and $g \geq 2$. Then, there exists a Stein surface W homotopic to the 2-sphere \mathbb{S}^2 which admits infinitely many Hamiltonian isotopy classes of embedded exact Lagrangian surfaces of genus g . In addition, W contains no embedded exact Lagrangian surfaces of genus h , $h \leq g - 1$.*

In Corollary 1.6, the Stein surface W for $g = 2$ can be constructed by attaching a Weinstein 2-handle to the standard symplectic 4-ball $(\mathbb{D}^4, \lambda_{\text{st}})$ along a max-tb Legendrian representative of the smooth knot $m(10_{145})$. The results we prove also allow us achieve $g = 1$ if we allow ourselves a bouquet of just two 2-spheres as the given homotopy type, instead of the 2-sphere \mathbb{S}^2 :

Corollary 1.7. *The Stein surface W obtained by attaching two Weinstein 2-handles along $\Lambda(\beta_{11}) \subset (\partial\mathbb{D}^4, \xi_{\text{st}})$, one per connected component, contains infinitely many Hamiltonian isotopy classes of embedded exact Lagrangian tori.*

Corollaries 1.6 and 1.7 are proven in Section 7. It remains an outstanding problem to construct a Legendrian knot with infinitely many distinct embedded Lagrangian 2-disk fillings (pairwise smoothly isotopic), or show no such knot exists.⁷

Organization. Here is an outline of the rest of the paper. In Section 2, we review some necessary background and formally describe the Legendrian links discussed in this introduction. The Floer-theoretical core of the article is developed in Sections 3, 4, and 5. In particular, Sections 3 and 4, jointly with Appendix A, develop a new combinatorial model for the maps between Legendrian contact DGAs with *integral coefficients* associated to a decomposable exact Lagrangian cobordism. We believe these results are of independent interest for 3-dimensional contact topology and Floer theory. We then apply these maps in Sections 6 to prove Theorem 1.1, and prove a number of corollaries and other ancillary results in Section 7.

Acknowledgements. We thank Tobias Ekholm, Honghao Gao, Eugene Gorsky, Linhui Shen, and Daping Weng for illuminating conversations. R. Casals is supported by the NSF grant DMS-1841913, the NSF CAREER grant DMS-1942363 and the Alfred P. Sloan Foundation. L. Ng is partially supported by the NSF grants DMS-1707652 and DMS-2003404. \square

⁷The case of a link with infinitely many planar Lagrangian fillings (pairwise smoothly isotopic) might already be an interesting start. In terms of Stein surfaces, the analogue of the knot case would be to construct a Stein surface homotopic to \mathbb{S}^2 with infinitely many Hamiltonian isotopy classes of pairwise smoothly isotopic Lagrangian 2-spheres.

In this section we describe the classes of Legendrian links $\Lambda \subset (\mathbb{R}^3, \xi_{\text{st}})$ and Legendrian loops that we study in this article. We begin in Section 2.1 with a review of Legendrian links and exact Lagrangian cobordisms, and then proceed in Sections 2.2 and 2.3 to describe the particular links of interest to us, which include the links in \mathcal{H} presented in the introduction. We conclude in Section 2.4 by describing the purple-box Legendrian loops that are a key ingredient in our constructions.

2.1. Legendrian links, exact Lagrangian cobordisms, and fillings. Here we briefly review the basic geometric terminology that we will need for this paper. There is now an extensive literature on exact Lagrangian cobordisms, including the papers cited in the introduction, to which we refer the reader for further details; specifically, the paper [EHK16] has a full exposition of the setting we will use here.

Rather than work with the contact manifold $(\mathbb{S}^3, \xi_{\text{st}})$ directly, it is convenient to remove a point and work in the contact manifold $(\mathbb{R}^3, \xi_{\text{st}})$, where ξ_{st} is the contact structure given by the kernel of the standard contact 1-form $\alpha_{\text{st}} := dz - y dx$ on \mathbb{R}^3 , endowed with Cartesian coordinates $(x, y, z) \in \mathbb{R}^3$. By definition, a link $\Lambda \subset (\mathbb{R}^3, \xi_{\text{st}})$ is *Legendrian* if it is everywhere tangent to ξ_{st} , or equivalently if $\alpha_{\text{st}}|_{\Lambda} = 0$; all Legendrian links in this paper are oriented.

As is customary, we will describe Legendrian links in $(\mathbb{R}^3, \xi_{\text{st}})$ by their front and Lagrangian projections. These are the images of the link under the projections $\Pi_{xz}, \Pi_{xy} : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ to the xz and xy planes, respectively. Given a Legendrian link Λ , the Reeb chords of Λ are integral curves of the Reeb vector field ∂_z with endpoints on Λ ; these correspond to the crossings of the Lagrangian projection $\Pi_{xy}(\Lambda)$. One numerical invariant associated to a Legendrian link Λ is the *Thurston–Bennequin number* $tb(\Lambda)$, which is the number of crossings of $\Pi_{xy}(\Lambda)$ counted with sign.

Example 2.1. The simplest Legendrian knot in \mathbb{R}^3 is the standard Legendrian unknot with $tb = -1$, which we will denote by U . The front projection $\Pi_{xz}(U)$ is a “flying saucer” with two cusps, while the Lagrangian projection $\Pi_{xy}(U)$ is a “figure eight” diagram with a single crossing; see the top left of Figure 5. \square

The *symplectization* of \mathbb{R}^3 is the 4-manifold $\mathbb{R}^4 = \mathbb{R}_t \times \mathbb{R}^3$ equipped with the exact symplectic form $\omega_{\text{st}} = d\lambda_{\text{st}}$ with $\lambda_{\text{st}} = e^t \alpha_{\text{st}}$. Note that this symplectic manifold is symplectomorphic to $(\mathbb{R}^4, d\lambda_0)$, where $\lambda_0 := \frac{1}{2}(x_1 dy_1 - y_1 dx_1 + x_2 dy_2 - y_2 dx_2)$, $(x_1, y_1, x_2, y_2) \in \mathbb{R}^4$, is the radial Liouville form in \mathbb{R}^4 . Given that $(\mathbb{R}^4, d\lambda_{\text{st}})$ is symplectomorphic to the Liouville completion of the standard symplectic Darboux ball $(\mathbb{D}^4, d\lambda_0)$, we will also write λ_{st} for λ_0 and denote by $(\mathbb{D}^4, \lambda_{\text{st}})$ the unique exact symplectic filling of $(\mathbb{S}^3, \xi_{\text{st}})$, with a radial primitive Liouville form.

We will be interested in *Lagrangian submanifolds* of $(\mathbb{R}^4, d\lambda_{\text{st}})$, which are surfaces $L \subset \mathbb{R}^4$ such that $\omega_{\text{st}}|_L = 0$. One class of Lagrangian submanifolds is given by cylinders over Legendrians: if $\Lambda \subset \mathbb{R}^3$ is Legendrian, then $\mathbb{R} \times \Lambda \subset \mathbb{R}^4$ is Lagrangian.

More generally, suppose that Λ_+, Λ_- are Legendrian links in \mathbb{R}^3 . A *Lagrangian cobordism* from Λ_- to Λ_+ is a Lagrangian $L \subset \mathbb{R}^4$ such that, for some $T > 0$,

$$L \cap ((-\infty, -T) \times \mathbb{R}^3) = (-\infty, -T) \times \Lambda_- \quad \text{and} \quad L \cap ((T, \infty) \times \mathbb{R}^3) = (T, \infty) \times \Lambda_+.$$

The Lagrangian cobordism L is *exact* if there is a function $f : L \rightarrow \mathbb{R}$ such that $\lambda_{\text{st}}|_L = df$ and f is constant on each of the ends $(-\infty, -T) \times \Lambda_-$ and $(T, \infty) \times \Lambda_+$ separately. All Lagrangian cobordisms considered in this paper will be oriented, embedded, and exact. In the special case where the negative end is empty, an exact Lagrangian cobordism from \emptyset to Λ is called a *filling* of Λ . See Figure 4.

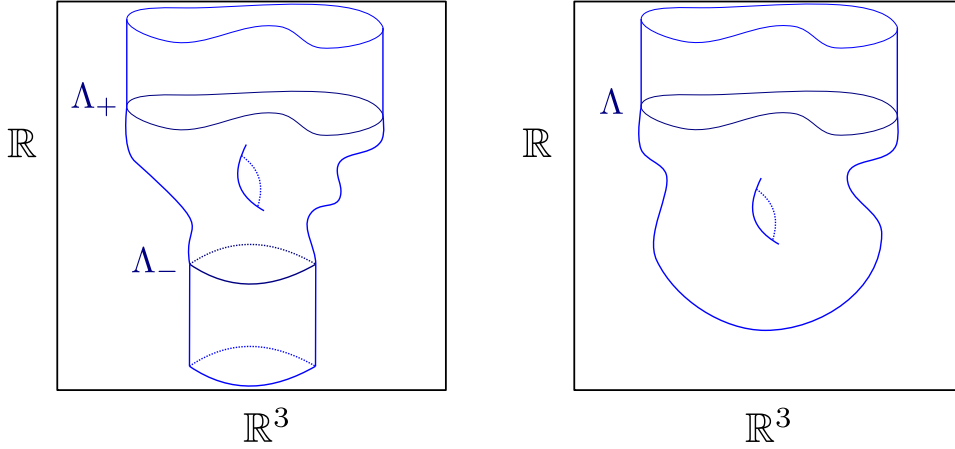


FIGURE 4. A Lagrangian cobordism from Λ_- to Λ_+ (left) and a filling of Λ (right).

We will be interested in Lagrangian cobordisms and fillings up to *exact Lagrangian isotopy*, which is an isotopy through exact Lagrangian cobordisms that fixes the two cylindrical ends (or the positive cylindrical end, in the case of fillings). In the setting of \mathbb{R}^4 , this is the same as a Hamiltonian isotopy, see e.g. [Oh15, Section 3.6], which is an isotopy through Hamiltonian diffeomorphisms fixing the two ends $(-\infty, -T) \times \mathbb{R}^3$ and $(T, \infty) \times \mathbb{R}^3$.

Remark 2.2. Associated to a Legendrian link in \mathbb{R}^3 or a Lagrangian surface in \mathbb{R}^4 is its *Maslov number*, which takes values in \mathbb{Z} . For a Lagrangian surface L , this is the greatest common divisor of the Maslov numbers of all closed loops in L , where the Maslov number of a loop in L is understood to be the Maslov number of the corresponding loop in the Lagrangian Grassmannian of \mathbb{R}^4 . For a Legendrian link Λ , the Maslov number is the Maslov number of the surface $\mathbb{R} \times \Lambda$. All Legendrians and Lagrangians that we consider in this paper will have Maslov number 0. \square

We will construct exact Lagrangian cobordisms out of key building blocks called *elementary cobordisms*, due to [EHK16]. There are three types of elementary cobordisms between Legendrian links, which we describe in turn.

(i) Isotopy cobordisms. If Λ_- and Λ_+ are Legendrian links that are related by a Legendrian isotopy Λ_t , then the trace of this isotopy (the union of $\{t\} \times \Lambda_t$ over all t) can be perturbed to an exact Lagrangian cobordism from Λ_- to Λ_+ , which we will call the *isotopy cobordism* associated to this isotopy. The isotopy represents a path in the space of Legendrian links from Λ_- to Λ_+ , and homotopic paths lead to isotopy cobordisms that are exact Lagrangian isotopic.

(ii) Minimum cobordisms. Let U denote a standard Legendrian unknot as in Example 2.1. By [EP96], U has a filling by a Lagrangian 2-disk, which is necessarily exact, and this filling is unique up to exact Lagrangian isotopy. Thus if Λ_- is any Legendrian link and Λ_+ is the split union of Λ_- and a standard unknot U , then there is an exact Lagrangian cobordism from Λ_- to Λ_+ given by the union of the filling of U and the cylinder $\mathbb{R} \times \Lambda_-$. This cobordism is called a *minimum cobordism* and corresponds topologically to the addition of a 0-handle.

(iii) Saddle cobordisms. Let Λ_+ be a Legendrian link. Reeb chords of Λ_+ correspond to crossings in the Lagrangian projection $\Pi_{xy}(\Lambda_+)$. A Reeb chord is called *contractible* if there is a Legendrian isotopy of Λ_+ inducing a planar isotopy of $\Pi_{xy}(\Lambda_+)$ and ending in a Legendrian where the height of the Reeb chord is arbitrarily small. Suppose that we have a contractible Reeb chord a of Λ_+ that corresponds to a positive crossing of $\Pi_{xy}(\Lambda_+)$ (in symplectic terms, the Conley–Zehnder index of a is even). One can modify the diagram $\Pi_{xy}(\Lambda_+)$ by replacing the corresponding crossing by its oriented resolution to produce the Lagrangian projection of another Legendrian link Λ_- ; see Figure 5. There is then an exact

Lagrangian cobordism from Λ_- to Λ_+ called a *saddle cobordism*. This is sometimes called a *pinch move* because of what it looks like in the front projection, and we will also sometimes refer to this as “resolving” the Reeb chord; it corresponds topologically to the addition of a 1-handle.




FIGURE 5. Two elementary Lagrangian cobordisms, depicted in terms of their xy projections: a minimum cobordism (left) and a saddle cobordism (right). The dotted arrows go from the bottom to the top of the cobordisms. The diagram on the top left is the standard Legendrian unknot U .

We can build more cobordisms out of elementary pieces through the operation of *concatenation*. Suppose that L_1 and L_2 are exact Lagrangian cobordisms that go from Λ_0 to Λ_1 and from Λ_1 to Λ_2 , respectively. We can remove the top cylinder of L_1 and the bottom cylinder of L_2 and glue the resulting Lagrangians along their common boundary Λ_1 to produce a new exact Lagrangian cobordism $L_1 \# L_2$ from Λ_0 to Λ_2 , the concatenation of L_1 and L_2 . An exact Lagrangian cobordism is *decomposable* if it is the concatenation of some number of elementary cobordisms. All of the cobordisms and fillings that we consider in this paper will be decomposable. Now that we have reviewed the basic geometric concepts and terminology, let us delve into the specific objects of interest with a view towards the new contributions of this manuscript.

2.2. Legendrian links associated to positive braids. We now describe the specific Legendrian links in $(\mathbb{R}^3, \xi_{\text{st}})$ that we will consider in this paper. These are a natural family of Legendrian links associated to positive braids, topologically given by the closures of these braids with one full negative twist.

Let Br_N denote the N -strand braid group, $N \in \mathbb{N}$. The standard presentation of Br_N is given by Artin generators $\sigma_1, \dots, \sigma_{N-1}$, where σ_i corresponds to a single positive crossing between strands i and $i+1$ of the braid, with relations $\sigma_i \sigma_{i+1} \sigma_i = \sigma_{i+1} \sigma_i \sigma_{i+1}$ for all i and $\sigma_i \sigma_j = \sigma_j \sigma_i$ for $|i-j| > 1$. Within Br_N , let Br_N^+ denote the monoid of positive braids; any element β of Br_N^+ can be written as a braid word

$$\beta := \prod_{j=0}^{l(\beta)} \sigma_{i_j}, \quad i_j \in [1, N-1],$$

where $l(\beta)$ is the length of $\beta \in \text{Br}_N^+$, equivalently its number of crossings.

Given a positive braid β , the *rainbow closure* of β is the Legendrian link whose front projection is given by drawing β horizontally and joining the left and right ends of β by a nested set of non-intersecting arcs with a single left and right cusp; see the top diagrams in Figure 7. Topologically this link is the 0-framed closure of β . As mentioned in the introduction, rainbow closures are the subject of several other papers on fillings of Legendrian links, including [CG21, CZ21, GSW20a, GSW20b].

We note that not all Legendrian links are (isotopic to) rainbow closures of positive braids. In particular, if Λ is the rainbow closure of a positive braid $\beta \in \text{Br}_N^+$, then Λ has Thurston–Bennequin number $tb(\Lambda) = l(\beta) - N$. If we write $g(\Lambda)$ for the Seifert genus of the topological link type of Λ , then there is an obvious Seifert surface for Λ whose Euler characteristic is $N - l(\beta)$. It follows that the Bennequin inequality $tb(\Lambda) \leq 2g(\Lambda) - 1$ must be sharp in this case, and in particular that Λ must maximize Thurston–Bennequin number within its topological type. Thus even if Λ represents a topological link that is a positive braid closure, it can only be a rainbow closure if it maximizes tb .

We will focus on another Legendrian link associated to a positive braid β , which we describe next and call the (-1) -closure of β . This is topologically the closure of β with a full negative twist, and is arguably more naturally associated to β than the rainbow closure, due to its connection to Legendrian satellites as described below. We remark that any rainbow closure of a positive braid is also (Legendrian isotopic to) the (-1) -closure of another braid, namely the concatenation of the original braid with a full positive twist.

There is a well-defined (up to isotopy) Legendrian link $\tilde{\Lambda}(\beta) \subset (J^1\mathbb{S}^1, \xi_{\text{st}})$ associated to β (cf. [EV18]). By definition, the Legendrian $\tilde{\Lambda}(\beta) \subset (J^1\mathbb{S}^1, \xi_{\text{st}})$ is the Legendrian link whose front in $\mathbb{S}^1 \times \mathbb{R}$ (image of the projection map $J^1\mathbb{S}^1 = T^*\mathbb{S}^1 \times \mathbb{R} \rightarrow \mathbb{S}^1 \times \mathbb{R}$) consists of the N horizontal strands $\mathbb{S}^1 \times \{j\}$, $j = 1, \dots, N$, where (positive) crossings are added left to right according to the braid word β . Figure 6 depicts $\tilde{\Lambda}(\beta)$ with an explicit example: the \mathbb{S}^1 -coordinate is horizontal, and the two vertical yellow walls are identified with each other.




FIGURE 6. The front for the Legendrian link $\tilde{\Lambda}(\beta) \subset (J^1\mathbb{S}^1, \xi_{\text{st}})$ associated to the braid word β (left). Explicit example of $\tilde{\Lambda}(\beta)$ associated to the braid word $\beta = \sigma_1\sigma_5\sigma_4\sigma_2\sigma_3\sigma_5\sigma_4\sigma_3\sigma_4\sigma_3\sigma_2\sigma_4\sigma_3\sigma_2^2\sigma_4\sigma_5\sigma_3\sigma_4\sigma_2\sigma_5\sigma_1\sigma_2 \in \text{Br}_6^+$ (right).

Given a Legendrian link $\tilde{\Lambda}(\beta) \subset (J^1\mathbb{S}^1, \xi_{\text{st}})$, we denote by $\Lambda(\beta) \subset (\mathbb{R}^3, \xi_{\text{st}})$ the Legendrian link obtained by satelliting $\tilde{\Lambda}(\beta)$ along the standard Legendrian unknot $U \subset \mathbb{R}^3$. To be precise, $(J^1\mathbb{S}^1, \xi_{\text{st}})$ is contactomorphic to a standard contact neighborhood $\mathcal{O}p(\Lambda_0) \subset (\mathbb{R}^3, \xi_{\text{st}})$ of $\Lambda_0 \subset (\mathbb{R}^3, \xi_{\text{st}})$, and $\Lambda(\beta)$ is the image of $\tilde{\Lambda}(\beta)$ under the resulting inclusion $J^1\mathbb{S}^1 \hookrightarrow \mathbb{R}^3$; this is a special case of the Legendrian satellite construction [NT04].

Figure 7 (bottom left) shows the front projection for the Legendrian link $\Lambda(\beta)$. The transition from $\Pi_{xz}(\Lambda)$ to $\Pi_{xy}(\Lambda)$ (“resolution”) can be performed as in [Ng03, Proposition 2.2] and it is also depicted in Figure 7, both for rainbow and (-1) -closures. In this Lagrangian projection, a combinatorial advantage is that Reeb chords for $\Lambda(\beta) \subset (\mathbb{R}^3, \xi_{\text{st}})$ are in bijection with the (positive) crossings of $\Pi_{xy}(\Lambda(\beta))$.

We will be interested in the Legendrian contact DGA and associated monodromy of $\Lambda(\beta)$, both of which can be combinatorially described via the xy projection of $\Lambda(\beta)$. Rather than use the xy projection shown in the bottom right of Figure 7, it will significantly simplify our computations to change $\Lambda(\beta)$ by a Legendrian isotopy to have a slightly different xy projection, as we describe next.




FIGURE 7. Two different Legendrian links Λ associated to a positive braid β , depicted in their front projection (left) and Lagrangian projection (right). The top row depicts the rainbow closure of β ; the bottom row is $\Lambda(\beta)$. The Lagrangian projections are obtained by resolving the corresponding front projections. The arrows indicate the orientation of the link.

Definition 2.3. Let β be a positive braid. Consider the link diagram in \mathbb{R}^2 given by the blackboard-framed satellite closure of β around the figure-eight unknot diagram $\Pi_{xy}(U)$, as depicted in the rightmost diagram of Figure 8. If this diagram is the Lagrangian projection of a Legendrian link, then we call this Legendrian link the (-1) -closure of β . \square

It is apparent that $\Lambda(\beta)$ and the (-1) -closure of β represent smoothly isotopic links, as they are both the (-1) -framed closures of the braid β . Furthermore, their Lagrangian projections are regularly homotopic: an isotopy between them is indicated in Figure 8. The first step in this isotopy is just a planar isotopy moving the $\binom{N}{2}$ negative crossings to the left of β to the top of the diagram. We then use a sequence of Reidemeister II and III moves to obtain the square-grid configuration of crossings shown in the blue box in Figure 8 (right). However, the smooth isotopy from the center diagram to the right diagram does not always represent a Legendrian isotopy—and in particular the right diagram does not even necessarily represent a Legendrian link—as we illustrate by an example.




FIGURE 8. An isotopy from the Lagrangian projection obtained by the resolution of a front projection (left) to the (-1) -closure (right). The red exclamation mark indicates that the smooth isotopy between the center and right diagrams does not necessarily represent a Legendrian isotopy, depending on the choice of β .

Example 2.4. Consider the Legendrian link $\Lambda(e)$ where e is the trivial 2-stranded braid $[\emptyset] \in \text{Br}_2^+$. Following the resolution procedure as in Figure 7, we find that the Lagrangian

projection of $\Lambda(e)$ is the exact Lagrangian $L \subset \mathbb{R}^2$ depicted in Figure 9.(i). In what follows, we use D. Sauvaget’s calculus [Sau04, Section II.2] for exact Lagrangian projections – see also [Lin16, Section 2] for an introduction. Let $A, B, C, P_1, P_2 \in \mathbb{R}^+$ be the areas of the bounded regions $\mathbb{R}^2 \setminus L$, as shown in Figure 9.(i); we may and do assume that we have $B < C$. The two area constraints for this projection read

$$P_1 = A + B, \quad P_2 = A + C.$$

In order to perform a Reidemeister III in the region with area B , we first empty the area in that region, leading to Figure 9.(ii), and the corresponding exactness constraints are satisfied:

$$P_1 = 0 + (A + B), \quad P_2 = (A + B) - (A + C).$$

The Reidemeister III move leads to Figure 9.(iii) and an additional Reidemeister II move, creating a canceling pair of crossings, to Figure 9.(iv). A second Reidemeister III move, which is admissible due to the zero area in its triangular region, yields Figure 9.(v). The area constraints are still satisfied, as they coincide with those in Figure 9.(ii). These moves concatenate to a Hamiltonian isotopy from Figure 9.(i) to Figure 9.(v), through exact Lagrangians. Now, we claim that the transition from Figure 8 (Center) to Figure 8 (right) cannot exist through exact Lagrangians: the resulting Lagrangian – shown in Figure 9 – is not an exact Lagrangian. This can be directly seen by the area constraints:

$$\alpha = \gamma + \delta + \varepsilon, \quad \alpha + \beta + \varepsilon = \gamma, \quad \alpha, \beta, \gamma, \delta, \varepsilon \in \mathbb{R}^+,$$

which imply $\delta + \beta + 2\varepsilon = 0$, contradicting positivity of the areas $\delta, \beta, \varepsilon \in \mathbb{R}^+$. Alternatively, it is rather immediate that the two curves in Figure 9.(vi) bound an immersed annulus, with positive area. Hence, the conclusion is that a constraint on β needs to be imposed, should we want to work with a Legendrian link through a Lagrangian projection of the form shown in Figure 8 (right). \square

We will want to consider braids where the isotopy in Figure 8 is legal. In the following definition, let $\pi : J^1\mathbb{S}^1 \rightarrow T^*\mathbb{S}^1$ denote the projection to the first factor, where $J^1\mathbb{S}^1 = T^*\mathbb{S}^1 \times \mathbb{R}_z$ with the standard contact form $dz - \lambda_{\text{st}}$.

Definition 2.5. Let $\beta \in \text{Br}_N^+$, $N \in \mathbb{N}$, and consider its smooth braid closure $c(\beta)$ in $\mathbb{S}^1 \times \mathbb{R} \cong T^*\mathbb{S}^1$, depicted as a (horizontal) link diagram. Then $\beta \in \text{Br}_N^+$ is said to be *admissible* if $c(\beta)$ is the Lagrangian projection of a Legendrian link $\Lambda \subset (J^1\mathbb{S}^1, \xi_{\text{st}})$: that is, if there exists a Legendrian link $\Lambda \subset (J^1\mathbb{S}^1, \xi_{\text{st}})$ such that $c(\beta) = \pi(\Lambda)$ as link diagrams, where crossings are taken into account. \square

As we now explain, if β is admissible, then the isotopy in Figure 8 is legal and in particular it makes sense to refer to the (-1) -closure of β .

Proposition 2.6. *Suppose that β is admissible. Then the diagram on the right of Figure 8 is the Lagrangian projection of a Legendrian link in \mathbb{R}^3 , and the sequence of moves in Figure 8 represents a Legendrian isotopy.*

Proof. Suppose β is admissible, and let Λ be the Legendrian link in $J^1\mathbb{S}^1$ whose Lagrangian projection is $c(\beta)$. If we cut Λ at a point in \mathbb{S}^1 then we obtain a Legendrian braid in $J^1([0, 1])$ in the terminology of [EV18]. By the classification of positive Legendrian braids [EV18, Theorem 3.4], this braid is Legendrian isotopic to the Legendrian braid whose *front* projection is β . It follows that this remains true when we satellite these braids around the standard Legendrian unknot $U \subset \mathbb{R}^3$. The satellite of the latter braid is $\Lambda(\beta)$ as defined in Section 2.2, which in the Lagrangian projection is the leftmost diagram in Figure 8. On the other hand, one directly sees (without passing to the front projection) that the Lagrangian projection of the satellite of Λ is the rightmost diagram in Figure 8. The result follows. \square




FIGURE 9. The sequence of Lagrangian projections discussed in Example 2.4. The transitions from (i) to (v) are all realized by Hamiltonian isotopies, preserving the exactness of the immersed Lagrangian. Item (vi) displays an example of a Lagrangian which is *not* exact. The consequence of (v) not being necessarily Legendrian isotopic to (vi) leads to Definition 2.5.

Example 2.4 shows that not every braid $\beta \in \text{Br}_N^+$ is admissible. Let us introduce a sufficiency criterion for a braid $\beta \in \text{Br}_N^+$ to be admissible. For that, let

$$\Delta_N = \prod_{i=1}^{N-1} \prod_{j=1}^{N-i} \sigma_j \in \text{Br}_N^+$$

denote the half-twist on N strands, i.e., the Garside element of the N -stranded braid group Br_N .

Proposition 2.7. *Any positive braid containing a half-twist is admissible, i.e. if β_1, β_2 are braids in Br_N^+ , then $\beta_1 \Delta_N \beta_2$ is admissible.*

Proof. Since admissibility depends only on the closure of the braid in the solid torus, we may move β_2 to the beginning of the braid; it thus suffices to show that if $\beta \in \text{Br}_N^+$ then $\beta \Delta_N$ is admissible. For this, consider the standard front for $\tilde{\Lambda}(\beta \Delta_N)$ in $\mathbb{S}^1 \times \mathbb{R}$ and deform it, scanning left-to-right, using the resolution procedure in [Ng03, Section 2.1]: see [Ng03, Figure 3] and Figure 10 (left). The procedure described in [Ng03] uses a front projection in $\mathbb{R}_q \times \mathbb{R}_z$, instead of $\mathbb{S}_q^1 \times \mathbb{R}_z$, but can still be used with this latter base $\mathbb{S}^1 \times \mathbb{R}$ by using the half-twist $\Delta_N \in \text{Br}_N^+$, which is part of the braid $\beta \Delta_N$ by hypothesis. Indeed, this is depicted in Figure 10 (left), where the half-twist is shown in the yellow box. The Lagrangian projection associated to this deformed front is depicted in Figure 10 (right), where the half-twist Δ_N now appears thanks to the crossings associated to the (green) Reeb chords that appear at the right-most part of the front in $\mathbb{S}_q^1 \times \mathbb{R}_z$. This concludes the statement. \square




FIGURE 10. Deforming the front projection for $\tilde{\Lambda}(\beta\Delta_N)$ in $\mathbb{S}^1 \times \mathbb{R}$ (left) so that the corresponding Lagrangian projection in $T^*\mathbb{S}^1$ is as shown on the right.

For future reference we note the following variant on Proposition 2.7. Each crossing in the Lagrangian projection in $T^*\mathbb{S}^1$ of a Legendrian link in $J^1\mathbb{S}^1$ corresponds to a Reeb chord of the link. A Reeb chord is called *contractible* if its height can be made arbitrarily small without changing the Lagrangian projection of the link (up to planar isotopy).

Proposition 2.8. *If β_1, β_2 are braids in Br_N^+ , then any crossing coming from β_1 or β_2 in the admissible braid $\beta_1\Delta_N\beta_2$ is contractible.*

Proof. First consider the special case where β_2 consists of a single crossing. We claim that this crossing is contractible. Indeed, a slight variant on the construction from Figure 10 involving swapping two of the strands in the yellow box in the front projection gives the desired contractible crossing: see Figure 11.




FIGURE 11. A variant on the argument from Proposition 2.7. Here two of the strands passing through the yellow Δ_N box in the front projection swap places, producing a contractible crossing in the Lagrangian projection: the red dot on the right, corresponding to the single crossing of β_2 .

In the general case, cut the closure of the braid $\beta_1\Delta_N\beta_2$ at the specified crossing. Push Δ_N to the end of the resulting braid by a sequence of Reidemeister III moves. From the above special case, we can realize the resulting braid as the Lagrangian projection of a Legendrian link in such a way that the distinguished crossing is contractible. Then push Δ_N back into its original position without disturbing a neighborhood of the contractible crossing; this is a braid isotopy and thus corresponds to a Legendrian isotopy by the classification of positive Legendrian braids [EV18]. \square

2.3. A class of Legendrian (-1) -closures. The Legendrian links that we use in this manuscript are particular examples of (-1) -closures, obtained by the following procedure. Let $w(\beta) \in S_N$ be the permutation given by the Coxeter projection $w : \text{Br}_N \rightarrow S_N$ of β onto the symmetric group, where the relations $\sigma_i^2 = 1$ are imposed for the Artin generators $i = 1, \dots, N - 1$. Suppose that the bijection $w(\beta) : [1, N] \rightarrow [1, N]$ has a fixed point i , for some $i \in [1, N]$. Then the Legendrian $\Lambda(\beta)$ contains a connected component $\Lambda(\beta)_i$

which is a standard Legendrian unknot. Since $\Lambda(\beta)_i \subset (\mathbb{R}^3, \xi_{\text{st}})$ is a Legendrian link, there exists a neighborhood $\mathcal{O}p(\Lambda(\beta)_i)$, disjoint from $\Lambda(\beta) \setminus \Lambda(\beta)_i$, which is contactomorphic to $\mathcal{O}p(\Lambda(\beta)_i) \cong (J^1\Lambda(\beta)_i, \xi_{\text{st}})$, where the contactomorphism sends $\Lambda(\beta)_i \subset \mathcal{O}p(\Lambda(\beta)_i)$ to the zero section $\Lambda(\beta)_i \subset (J^1\Lambda(\beta)_i, \xi_{\text{st}})$.

Now, let $\gamma \in \text{Br}_M^+$ be a positive M -stranded braid, $M \in \mathbb{N}$. Let us denote by $\Lambda(\gamma)_i \subset \mathcal{O}p(\Lambda(\beta)_i)$ the Legendrian link obtained by satelliting $\tilde{\Lambda}(\gamma) \subset (J^1\mathbb{S}^1, \xi_{\text{st}})$ along the standard Legendrian unknot $\Lambda(\beta)_i \subset \mathcal{O}p(\Lambda(\beta)_i)$.

Definition 2.9. Let $\beta \in \text{Br}_N^+$ be such that $i \in [1, N]$ is a fixed point of $w(\beta)$, and $\gamma \in \text{Br}_M^+$ be an M -stranded braid, $M \in \mathbb{N}$. The Legendrian link $\Lambda(\beta, i; \gamma) \subset (\mathbb{R}^3, \xi_{\text{st}})$ is the Legendrian link $(\Lambda(\beta) \setminus \Lambda(\beta)_i) \cup \Lambda(\gamma)_i \subset (\mathbb{R}^3, \xi_{\text{st}})$, where $\Lambda(\gamma)_i \subset \mathcal{O}p(\Lambda(\beta)_i)$ is embedded in an arbitrarily but fixed neighborhood of the component $\Lambda(\beta)_i$. Colloquially, $\Lambda(\beta, i; \gamma)$ is the result of satelliting the braid γ around the component of the Legendrian link $\Lambda(\beta)$ labeled by i . \square

The Legendrian links in Theorem 1.1 are of the form $\Lambda(\beta, i; \gamma)$ for $\gamma \in \text{Br}_2^+$ and $\beta \in \text{Br}_N^+$, where $N = 2, 3$. For instance, the Legendrian links Λ_n come from setting $\beta = \sigma_1^6$ and $\gamma = \sigma_1^n$ with $N = M = 2$:

$$\Lambda_n \cong \Lambda(\sigma_1^6, 1; \sigma_1^n).$$

Similarly, the Legendrian links $\Lambda(\tilde{D}_n)$, $n \geq 4$, come from setting $\beta = (\sigma_1\sigma_2\sigma_2\sigma_1)^2\sigma_2^2$ and $\gamma = \sigma_1^{n-2}$ with $N = 3, M = 2$:

$$\Lambda(\tilde{D}_n) \cong \Lambda((\sigma_1\sigma_2\sigma_2\sigma_1)^2\sigma_1^2, 1; \sigma_1^{n-2}).$$

Remark 2.10. As noted in the introduction, the Legendrian links $\Lambda(\tilde{D}_n)$, $n \geq 4$, are also the rainbow closures of the positive braids

$$\eta_n = (\sigma_2\sigma_1\sigma_3\sigma_2\sigma_2\sigma_3\sigma_1\sigma_2)\sigma_1^{n-4}, \quad n \geq 4, \quad \eta_n \in \text{Br}_4^+.$$

The brick diagram [Rud92, BLL18] associated to this positive braid word η_n coincides with the CoxeterDynkin diagrams \tilde{D}_n associated to the affine Coxeter group of D -type. This affine Coxeter diagram also arises from two natural constructions starting with η_n . First, the quiver associated to the positive braid η_n , according to the algorithm in [BFZ05], and second, as the diagram for the intersection form associated to a set of (distinguished) generators in the first homology group of a minimal-genus Seifert surface associated to the link given by η_n [Mis17, BLL18]. In addition, the augmentation variety associated to $\Lambda(\tilde{D}_n)$ admits a cluster structure of \tilde{D}_n -type. These reasons lead us to the notation $\Lambda(\tilde{D}_n)$ and referring to these braids as the (maximal-tb) affine D_n -Legendrian links. \square

Definition 2.9 is rather direct diagrammatically. Indeed, given the front diagram for $\Lambda(\beta) \subset (\mathbb{R}^3, \xi_{\text{st}})$ shown in Figure 7 (left), a front diagram for $\Lambda(\beta, i; \gamma) \subset (\mathbb{R}^3, \xi_{\text{st}})$ is obtained by taking the M -copy Reeb push-off of the i -th component of $\Lambda(\beta)$, corresponding to the i -th strand in Br_N^+ , and inserting the front diagram for $\tilde{\Lambda}(\gamma)$. This is shown in Figure 12.




FIGURE 12. Front projections for the Legendrian links $\Lambda(\beta) \subset (\mathbb{R}^3, \xi_{\text{st}})$ (left) and $\Lambda(\beta, i; \gamma) \subset (\mathbb{R}^3, \xi_{\text{st}})$ (right).

Similarly, this construction is depicted in the Lagrangian projection in Figure 13.




FIGURE 13. Lagrangian projections for the Legendrian links $\Lambda(\beta) \subset (\mathbb{R}^3, \xi_{\text{st}})$ (left) and $\Lambda(\beta, i; \gamma) \subset (\mathbb{R}^3, \xi_{\text{st}})$ (right). These are the Lagrangian projections that we use in order to compute the Legendrian contact DGA.

The crucial property of the Legendrian links $\Lambda(\beta, i; \gamma) \subset (\mathbb{R}^3, \xi_{\text{st}})$ is the existence of a specific contact isotopy $\varphi_t : (\mathbb{R}^3, \xi_{\text{st}}) \rightarrow (\mathbb{R}^3, \xi_{\text{st}})$, $t \in [0, 1]$, such that $\varphi_1(\Lambda(\beta, i; \gamma)) = \Lambda(\beta, i; \gamma)$ and $\varphi_t|_{\mathbb{R}^3 \setminus \mathcal{O}p(\Lambda(\beta)_i)} = \text{Id}$ for all $t \in [0, 1]$, as we now explain.

2.4. The purple-box Legendrian loop. Let $\beta \in \text{Br}_N^+$, $\gamma \in \text{Br}_M^+$ and consider the Legendrian link $\Lambda(\beta, i; \gamma) \subset (\mathbb{R}^3, \xi_{\text{st}})$. We construct a Legendrian loop $\vartheta : \mathbb{S}^1 \rightarrow \mathcal{L}(\Lambda(\beta, i; \gamma))$ based at $\Lambda(\beta, i; \gamma)$, whose action on the Legendrian contact DGA of $\Lambda(\beta, i; \gamma)$ will be studied in Section 5, and subsequently lead to Theorem 1.1. Intuitively, the Legendrian loop ϑ will fix the components of the Legendrian link $\Lambda(\beta)$ which do not belong to the satellite $\Lambda(\gamma) \subset \Lambda(\beta)$, and induce a rotation of $\Lambda(\gamma)$ corresponding to one full revolution of the \mathbb{S}^1 direction in $J^1\mathbb{S}^1$. Let us provide the details for its rigorous description.

Consider the component $\Lambda(\beta)_i \subset \Lambda(\beta)$ with a standard neighborhood $\mathcal{O}p(\Lambda(\beta)_i)$ and the Legendrian link $\tilde{\Lambda}(\gamma) \subset \mathcal{O}p(\Lambda(\beta)_i)$. Fix a contactomorphism

$$\mathcal{O}p(\Lambda(\beta)_i) \cong (J^1\mathbb{S}_\theta^1, \ker(dz - p_\theta d\theta)),$$

where $(J^1\mathbb{S}_\theta^1, \ker(dz - p_\theta d\theta))$ is the 1-jet space with coordinates $(\theta, p_\theta) \in T^*\mathbb{S}^1$, $z \in \mathbb{R}$. Fix the standard round metric in \mathbb{S}^1 , and choose $R \in \mathbb{R}^+$ such that $\tilde{\Lambda}(\gamma) \subset B_R$, where $B_R = \mathbb{D}_R(T^*\mathbb{S}^1) \times [-R, R] \subset T^*\mathbb{S}^1 \times \mathbb{R}$, with $\mathbb{D}_R(T^*\mathbb{S}^1)$ being the radius R (open) disk bundle.

Now, consider the Hamiltonian $p_\theta : J^1\mathbb{S}_\theta^1 \rightarrow \mathbb{R}$ and its associated contact vector field $X_{p_\theta} = -\partial_\theta$. Let $\varepsilon \in \mathbb{R}^+$, and choose a smooth cut-off function $\chi : J^1\mathbb{S}^1 \rightarrow \mathbb{R}$ such that

$$\chi|_{B_{R+\varepsilon}} \equiv 1, \quad \chi|_{B_{R+2\varepsilon} \setminus B_{R+\varepsilon}} \equiv 0.$$

The contact vector field $X_{\tilde{\vartheta}}$ associated to the Hamiltonian $\chi \cdot p_\theta : J^1\mathbb{S}^1 \rightarrow \mathbb{R}$ restricts to $-\partial_\theta$ in the tube B_R containing $\tilde{\Lambda}(\gamma)$, and it vanishes away from B_R . The contact flow of $X_{\tilde{\vartheta}}$ yields a compactly supported contact isotopy $\tilde{\Theta}_t : (J^1\mathbb{S}^1, \xi_{\text{st}}) \rightarrow (J^1\mathbb{S}^1, \xi_{\text{st}})$, which we parametrize such that $t = 1$ is the smallest $t \in \mathbb{R}^+$ with $\tilde{\Theta}_t(\tilde{\Lambda}(\gamma)) = \tilde{\Lambda}(\gamma)$ pointwise.

Definition 2.11. Let $\beta \in \text{Br}_N^+$ be such that $i \in [1, N]$ is a fixed point of $w(\beta)$, and let $\gamma \in \text{Br}_M^+$ be an M -stranded braid, $M \in \mathbb{N}$. The Θ_t -contact isotopy associated to $\Lambda(\beta, i; \gamma) \subset (\mathbb{R}^3, \xi_{\text{st}})$, $t \in [0, 1]$, is the compactly supported isotopy obtained by extending the compactly supported contact isotopy $\tilde{\Theta}_t : \mathcal{O}p(\Lambda(\beta)_i) \rightarrow \mathcal{O}p(\Lambda(\beta)_i)$, $t \in [0, 1]$, by the identity map on the complement of $\mathcal{O}p(\Lambda(\beta)_i)$. A Legendrian loop $\vartheta : \mathbb{S}^1 \rightarrow \mathcal{L}(\Lambda(\beta, i; \gamma))$ is said to be a ϑ -loop if it is obtained as $\Theta_t(\Lambda(\beta, i; \gamma))$, $t \in [0, 1]$, for a Θ_t -contact isotopy associated to $\Lambda(\beta, i; \gamma) \subset (\mathbb{R}^3, \xi_{\text{st}})$. \square

We will also call the Legendrian ϑ -loops in Definition 2.11 *purple-box Legendrian loops*, as they are obtained by moving the purple box which contains the braid γ clockwise around until it comes back to itself. Figure 14 (left) provides a schematic picture of such a ϑ -loop.




FIGURE 14. Left: a Legendrian ϑ -loop for the Legendrian link $\Lambda(\beta, i; \gamma) \subset (\mathbb{R}^3, \xi_{\text{st}})$, where the purple γ -box moves around clockwise around the β -box and comes back to itself using the upper strands). Right: the local move, consisting of a sequence of $l(\gamma)$ Reidemeister III moves, which we use in order to push the purple γ -box, right to left, through the β -box.

From a computational viewpoint, it is important to stress that a Legendrian ϑ -loop can be described in the Lagrangian projection strictly in terms of Reidemeister III moves and planar isotopies.⁸ In precise terms, a Legendrian ϑ -loop consists of two pieces:

- (i) Transferring the purple γ -box through the β -box, through a sequence of Reidemeister III moves. Indeed, it suffices to notice that moving the purple γ -box through one strand is achieved by $l(\gamma)$ consecutive Reidemeister III moves, one per each crossing of γ . This local move, past one strand, is shown in Figure 14 (right). Thus, the purple γ -box can be pushed through the β -box, right to left, by performing $l(\beta) \cdot l(\gamma)$ Reidemeister III moves.
- (ii) Moving the purple γ -box from the left of the β -box to its right *using the upper strands*. This is achieved by a planar isotopy, which moves the purple γ -box up and to the right (leaving the β -box beneath and passing above it), and then applying $N^2 \cdot l(\gamma)$ Reidemeister III moves to make the purple γ -box go around the pig-tailed loop until it returns to its initial position.

Hence, using a total of $l(\gamma) \cdot (N^2 + l(\beta))$ Reidemeister III moves in the Lagrangian projection, we can realize the Legendrian ϑ -loops in Definition 2.11.

Remark 2.12. Legendrian ϑ -loops can be considered as elements in $\pi_1(\mathcal{L}(\Lambda(\beta, i; \gamma)))$, or we can graph them in the symplectization as Lagrangian self-concordances $L_\vartheta \subset (\mathbb{R} \times \mathbb{R}^3, \lambda_{\text{st}})$ from the Legendrian link $\Lambda(\beta, i; \gamma) \subset (\mathbb{R}^3, \xi_{\text{st}})$ to itself. Most interestingly, given an exact Lagrangian filling $L \subset (\mathbb{D}^4, \lambda_{\text{st}})$ of $\Lambda(\beta, i; \gamma) \subset (\mathbb{S}^3, \xi_{\text{st}})$, we can concatenate L with L_ϑ , at the convex end of L and the concave end of L_ϑ . One may ask whether concatenating Lagrangian fillings with L_ϑ yields new Lagrangian fillings not Hamiltonian isotopic to L . Theorem 1.1 shows that there are Legendrian links where concatenating certain Lagrangian fillings with k consecutive copies of L_ϑ yields (infinitely many) pairwise distinct Lagrangian fillings, for different values of $k \in \mathbb{N}$. \square

Example 2.13. Legendrian ϑ -loops behave differently depending on the choice of braids $\beta \in \text{Br}_N^+$ and $\gamma \in \text{Br}_M^+$. For example, if $\gamma \in \text{Br}_1^+$ is the trivial 1-stranded braid, then the ϑ -loop is constant on the entire link $\Lambda(\beta, i; \gamma) \subset (\mathbb{R}^3, \xi_{\text{st}})$, regardless of the choice of

⁸One could instead use the resolution of the front projection as in Figure 7, and similarly push the purple γ -box around the front projection; this is e.g. what Kálmán does in [Kál05]. However, this version of the isotopy requires the use of both Reidemeister III and II moves. Our setup does not require Reidemeister II moves and this consequently simplifies our computations with the Legendrian contact DGA.

$\beta \in \text{Br}_N^+$. On the other hand, if we choose the braid β to be 1-stranded and the purple box $\gamma = \sigma_1^{n+2} \in \text{Br}_2^+$ to be 2-stranded, then we recover Kálmán’s Legendrian loop of $(2, n)$ -torus links [Kál05]. In this case, [Kál05, Theorem 1.3] shows that the action of the ϑ -loop on the degree-0 Legendrian contact homology of $\Lambda(\beta, i; \gamma)$ is nontrivial but of finite order. See Section 5.3 for further discussion of the Kálmán loop. \square

3. LEGENDRIAN CONTACT DGAs AND COBORDISM MAPS

In this section, we review the definition of the Legendrian contact DGA, with particular attention paid to integer and group-ring coefficients and the role of spin structures. We then proceed to discuss maps between DGAs induced by exact Lagrangian cobordisms, including exact Lagrangian fillings. There is now a reasonably large literature about these cobordism maps, beginning with work of Ekholm, Honda, and Kálmán [EHK16] defining the maps over \mathbb{Z}_2 ; we will need to compute a lift of these maps to \mathbb{Z} , which abstractly exists by work of Karlsson [Kar17, Kar20]. In this section we will present a framework that will allow us to perform explicit combinatorial computations of the cobordism maps over \mathbb{Z} , building them out of maps corresponding to particular elementary cobordisms. The maps for these elementary cobordisms are then presented in the following section, Section 4.

3.1. The Legendrian contact DGA. The Legendrian contact DGA, also known as the Chekanov–Eliashberg DGA, has been well-studied in the literature, especially in the setting of $(\mathbb{R}^3, \xi_{\text{st}})$. For the definition of the DGA in this setting, we refer the reader e.g. to [Che02] for the original definition over \mathbb{Z}_2 , [ENS02] for the definition over $\mathbb{Z}[t^{\pm 1}]$ (see also the survey [EN18]), and [NR13, NRS⁺20] for an upgraded definition with multiple base points. Here we will briefly review the definition that we will use, with \mathbb{Z} coefficients and multiple base points.

Let Λ be an oriented Legendrian link in $(\mathbb{R}^3, \xi_{\text{st}})$ equipped with a number of base points, such that there is at least one base point on each component. We will assume that Λ is sufficiently generic that the xy projection $\Pi_{xy}(\Lambda)$ in \mathbb{R}^2 is immersed with only transverse double point singularities, and no base point lies at one of these double points. We label the crossings of $\Pi_{xy}(\Lambda)$, which correspond to Reeb chords of Λ , as a_1, \dots, a_r , and decorate each base point with a monomial of the form $\pm s_i^{\pm 1}$. Let $\{s_1, \dots, s_q\}$ be the collection of indeterminates that appear in the labeling of the base points. To this decorated oriented Legendrian link Λ , we can associate the Legendrian contact DGA $(\mathcal{A}_\Lambda, \partial)$, as follows.

Generators. The algebra \mathcal{A}_Λ is the unital tensor algebra over the coefficient ring $\mathbb{Z}[s_1^{\pm 1}, \dots, s_q^{\pm 1}]$ generated by a_1, \dots, a_r . (One can lift this to the “fully noncommutative” algebra where the coefficients $s_i^{\pm 1}$ do not commute with Reeb chords a_i , and in our computations we will sometimes order our monomials accordingly. However, for the purposes of this paper, we will always assume that coefficients and Reeb chords commute.)

Grading. We assume for simplicity that each component of Λ has rotation number 0, which will be the case for the Legendrian links we study. The algebra \mathcal{A}_Λ is then graded over the integers \mathbb{Z} ; if Λ has a single component, then this grading is well-defined, while if Λ has multiple components, the grading depends on some additional choices. We will fix the grading by choosing a collection of distinguished base points, one on each component, such that the oriented tangent vectors to $\Pi_{xy}(\Lambda)$ at these points are all parallel in \mathbb{R}^2 . Label these base points by t_1, \dots, t_m , where m is the number of components of Λ and the base point t_j is on the j -th component. Consider a Reeb chord $a \in \mathcal{A}_\Lambda$ that ends on component $r(a)$ and begins on component $c(a)$; we define a capping path γ_a along Λ to be the concatenation of a path from the beginning point (undercrossing) of a to $t_{c(a)}$, and a path from $t_{r(a)}$ to the ending point of a , following the orientation of Λ for both paths. As we traverse γ_a , the unit tangent vector to $\Pi_{xy}(\gamma_a)$ changes continuously from the tangent vector to the undercrossing at a to

the tangent vector to the overcrossing; let $r(\gamma_a) \in \mathbb{R}$ denote the number of counter-clockwise revolutions around \mathbb{S}^1 that the tangent vector makes during this process, and note that $r(\gamma_a) \notin \frac{1}{2}\mathbb{Z}$ because of transversality. Then the grading of a is defined to be $-[2r(\gamma_a)] \in \mathbb{Z}$. We also place all the marked point monomials s_i in grading 0, which completes the grading of \mathcal{A}_Λ .

Differential. In order to set up the differential ∂ on \mathcal{A}_Λ , we first decorate the four quadrants at each crossing of $\Pi_{xy}(\Lambda)$ by two signs, a Reeb sign and an orientation sign. At each crossing, two opposite quadrants have Reeb sign $+$ and the others have Reeb sign $-$, while the orientation signs depend on whether the crossing is positive (even degree) or negative (odd degree): for positive crossings, two quadrants have orientation sign $+$ and two have $-$, while for negative crossings, all four quadrants have orientation sign $+$. See Figure 15.




FIGURE 15. In the top row, the Reeb signs (left diagram) and orientation signs (two right diagrams) at a crossing. Quadrants that have $-$ orientation sign are shaded, while all other quadrants have $+$ orientation sign. In the bottom row, two examples of disks in $\Delta(a)$. Both disks have $\text{sgn} = +1$ (on the right, the corner with negative orientation sign cancels the $-$ in $-s_1^{-1}$) and they contribute $+1$ and $+s_2^{-1}a_2a_1s_1^{-1}a_3$, respectively, to $\partial(a)$.

The differential now counts immersions of a disk D^2 with boundary punctures to \mathbb{R}^2 , mapping the boundary of D^2 to $\Pi_{xy}(\Lambda)$, such that a neighborhood of each boundary puncture is mapped to one of the four quadrants at a crossing of $\Pi_{xy}(\Lambda)$. We call such a disk an *immersed disk* for short; each corner of an immersed disk is a *positive (+) corner* or a *negative (-) corner* depending on the Reeb sign of the quadrant. For a Reeb chord a , define $\Delta(a)$ to be the set of immersed disks (up to reparametrization) with a single $+$ corner at a and no other $+$ corners. To any such disk $\Delta \in \Delta(a)$, we can define two quantities. One is the sign $\text{sgn}(\Delta) \in \{\pm 1\}$, given by the product of the orientation signs over all corners of Δ , multiplied by the signs of any base points traversed by the boundary of the disk ($+1$ for any base point labeled by $s_i^{\pm 1}$ and -1 for any base point labeled by $-s_i^{\pm 1}$). The other is the word $w(\Delta) \in \mathcal{A}_\Lambda$, which is the product, in order, of the Reeb chords at the $-$ corners and the base points that are encountered as we traverse the boundary of the disk counterclockwise, beginning and ending at the corner at a . A base point labeled by $\pm s_i^{\pm 1}$ contributes $s_i^{\pm 1}$ if it is traversed along the orientation of Λ and $s_i^{\mp 1}$ if it is traversed oppositely. The differential $\partial(a)$ is now defined to be:

$$\partial(a) := \sum_{\Delta \in \Delta(a)} \text{sgn}(\Delta)w(\Delta).$$

See Figure 15 for an example.

Remark 3.1 (multiple base points). In order to count augmentations over \mathbb{Z} , it is important that each component of Λ have at least one base point. Adding extra base points beyond one per component changes the DGA in a simple way. First note that moving a base point labeled $\pm s_i^{\pm 1}$ along Λ and through a crossing a has the effect of replacing a by $(\pm s_i^{\pm 1})^{\pm 1} a$: that is, the algebra is the same before and after the move, and the differential changes by conjugation by the automorphism that sends a to $(\pm s_i^{\pm 1})^{\pm 1} a$ and fixes all other generators. Thus if we have multiple base points on a single component, then up to a $\mathbb{Z}[s_1^{\pm 1}, \dots, s_q^{\pm 1}]$ -algebra isomorphism of the DGA, we can assume that all of the base points lie on the same segment of $\Pi_{xy}(\Lambda)$. In this case we can replace the multiple base points by a single base point labeled by the product of their labels, and the differential is unchanged. \square

Remark 3.2 (dependence on spin structure). In the differential over \mathbb{Z} of the Legendrian contact DGA $(\mathcal{A}_\Lambda, \partial)$ of a link Λ , the signs depend on a choice of spin structure on Λ , as laid out by the construction of Ekholm, Etnyre, and Sullivan [EES05]. For each \mathbb{S}^1 connected component of the Legendrian link Λ , there are two spin structures: the *Lie group* spin structure, induced by the fact that the 1-sphere \mathbb{S}^1 is a Lie group, and the *null-cobordant* spin structure, induced by the fact that \mathbb{S}^1 bounds a 2-disk D^2 and we can restrict the unique spin structure on D^2 to the boundary \mathbb{S}^1 . Here we review the discussion in [EES05] about how the choice of spin structure affects the differential ∂ in $(\mathcal{A}_\Lambda, \partial)$.

Choose one base point on each of the m components of Λ , so that \mathcal{A}_Λ is an algebra over $R = \mathbb{Z}[t_1^{\pm 1}, \dots, t_m^{\pm 1}]$, and write ∂^{comb} for the combinatorial differential on \mathcal{A}_Λ as defined above. The set of spin structures on Λ is an affine space based on $H_1(\Lambda, \mathbb{Z}_2) \cong \mathbb{Z}_2^m$; of interest to us will be two spin structures differing by $(1, \dots, 1)$, given by choosing the Lie group spin structure or the null-cobordant spin structure on *all* components of Λ . We will write ∂^{Lie} and ∂^{NC} for the geometric differentials on \mathcal{A}_Λ corresponding to these two spin structures.⁹ The two differentials ∂^{Lie} and ∂^{NC} depend on a number of auxiliary choices, including capping operators for Reeb chords—see Section 3.4 below for further discussion—but up to R -algebra isomorphism, $(\mathcal{A}_\Lambda, \partial^{\text{Lie}})$ and $(\mathcal{A}_\Lambda, \partial^{\text{NC}})$ are well-defined.

The combinatorial differential ∂^{comb} comes from the Lie group spin structure on Λ . To be precise, in [EES05, Theorem 4.32] it is shown that one can make choices so that ∂^{Lie} agrees with our definition of ∂^{comb} with signs as in Figure 15, except that for positive crossings (the left diagram on the top right of Figure 15), the opposite two quadrants are shaded.¹⁰ This change of shading corresponds to the R -algebra isomorphism of \mathcal{A}_Λ sending each Reeb chord a to $-a$ for even-graded Reeb chords and $+a$ for odd-graded Reeb chords, and so this isomorphism sends $(\mathcal{A}_\Lambda, \partial^{\text{comb}}) \xrightarrow{\cong} (\mathcal{A}_\Lambda, \partial^{\text{Lie}})$.

For cobordisms, the null-cobordant spin structure is more natural than the Lie group spin structure. To compute ∂^{NC} , we can appeal to [EES05, Theorem 4.29] (see also Remark 4.35 from the same paper), which implies that changing the spin structure by $(c_1, \dots, c_m) \in \mathbb{Z}_2^m$ has the effect of replacing t_i by $(-1)^{c_i} t_i$ for $i = 1, \dots, m$. In particular, define the \mathbb{Z} -algebra isomorphism $\phi : \mathcal{A}_\Lambda \rightarrow \mathcal{A}_\Lambda$ by $\phi(a) = a$ for all Reeb chords a and $\phi(t_i) = -t_i$ for all i ; then

$$\phi : (\mathcal{A}_\Lambda, \partial^{\text{Lie}}) \xrightarrow{\cong} (\mathcal{A}_\Lambda, \partial^{\text{NC}}).$$

More generally, suppose that we have multiple base points on each component of Λ as in Remark 3.1, each decorated by a monomial of the form $s_i^{\pm 1}$. Then, since no base point introduces a sign, the resulting combinatorial DGA $(\mathcal{A}_\Lambda, \partial^{\text{comb}})$ over $\mathbb{Z}[s_1^{\pm 1}, \dots, s_q^{\pm 1}]$ has signs corresponding to the Lie group spin structure. Now suppose that \mathcal{S} is any subset of these base points. If we replace the decoration $s_i^{\pm 1}$ of each base point in \mathcal{S} by $-s_i^{\pm 1}$, we obtain a new differential $\partial^{\mathcal{S}}$ on \mathcal{A}_Λ . Then $\partial^{\mathcal{S}}$ gives the differential corresponding to the spin

⁹The superscript ∂^{NC} stands for Null-Cobordant.

¹⁰In fact [EES05, Theorem 4.32] presents two choices of signs for ∂^{Lie} , of which we are describing one; however, it was subsequently proven in [Ng10] that the two choices lead to isomorphic DGAs.

structure that differs from the Lie group spin structure by $(c_1, \dots, c_m) \in \mathbb{Z}_2^m$, where c_i is the number of base points in \mathcal{S} that lie on component i . In particular, if \mathcal{S} has an odd number of points on each component¹¹, then we have an isomorphism of DGAs over $\mathbb{Z}[s_1^{\pm 1}, \dots, s_q^{\pm 1}]$:

$$(\mathcal{A}_\Lambda, \partial^{\mathcal{S}}) \cong (\mathcal{A}_\Lambda, \partial^{\text{NC}}).$$

□

3.2. Link automorphisms. In the case where Λ is a multi-component Legendrian link, rather than a knot, there is a structure on the Legendrian contact DGA of Λ that is hidden in the knot case. This is the “link grading” first introduced by K. Mishachev [Mis03], which essentially gives the DGA the structure of a path algebra (the “composable algebra”) on a graph whose vertices are components of Λ and whose edges are Reeb chords of Λ . This structure leads to a family of automorphisms of the DGA of the Legendrian link Λ , which we call *link automorphisms*. These will feature in our discussion at various points, and we discuss them now in detail.

Let $\Lambda = \Lambda_1 \cup \dots \cup \Lambda_m$ be an m -component Legendrian link. For any Reeb chord a of Λ , define $r(a), c(a) \in \{1, \dots, m\}$ to be the number of the component containing the endpoint (for $r(a)$) or beginning point (for $c(a)$) of a . The key observation of Mishachev is the following: in the DGA for Λ , any term in the differential ∂a of a Reeb chord a must be of the form $a_{i_1} \cdots a_{i_k}$, where $r(a) = r(a_{i_1}), c(a_{i_1}) = r(a_{i_2}), \dots, c(a_{i_{k-1}}) = r(a_{i_k}), c(a_{i_k}) = c(a)$. This motivates the following definition.

Definition 3.3. Let Λ be an m -component Legendrian link and $(\mathcal{A}_\Lambda, \partial)$ its DGA. A *link automorphism* of Λ is an algebra automorphism $\Omega : \mathcal{A}_\Lambda \rightarrow \mathcal{A}_\Lambda$ of the following form: there exist units u_1, \dots, u_m in the coefficient ring of \mathcal{A}_Λ such that for all Reeb chords a ,

$$\Omega(a) = u_{r(a)} u_{c(a)}^{-1} a.$$

□

The following is an immediate consequence of Mishachev’s observation.

Proposition 3.4. *Let Λ be a Legendrian link. Any link automorphism $\Omega : \mathcal{A}_\Lambda \rightarrow \mathcal{A}_\Lambda$ is a chain map of the Legendrian DGA $(\mathcal{A}_\Lambda, \partial)$.* □

In addition, Mishachev’s link grading structure is preserved by Legendrian isotopy, as can be checked by keeping track of components in the DGA chain maps induced by Legendrian isotopy. See [Mis03], and see Section 4.1 below for explicit formulas for these chain maps. As a consequence, link automorphisms persist under Legendrian isotopy:

Proposition 3.5. *Suppose Λ and Λ' are Legendrian isotopic links with respective DGAs $(\mathcal{A}_\Lambda, \partial)$ and $(\mathcal{A}_{\Lambda'}, \partial)$. Suppose that $\Psi : (\mathcal{A}_\Lambda, \partial) \rightarrow (\mathcal{A}_{\Lambda'}, \partial)$ is the DGA map induced by a Legendrian isotopy. If $\Omega : \mathcal{A}_\Lambda \rightarrow \mathcal{A}_\Lambda$ is a link automorphism of Λ , then there is a corresponding link automorphism $\Omega' : \mathcal{A}'_\Lambda \rightarrow \mathcal{A}'_\Lambda$ of Λ' such that $\Omega' \circ \Psi = \Psi \circ \Omega$.*

Proof. The numbering of the components of the Legendrian link Λ induces a corresponding numbering of the components of the link Λ' . If Ω is defined by $\Omega(a) = u_{r(a)} u_{c(a)}^{-1} a$ for some (u_1, \dots, u_m) , then we define Ω' in the same way: $\Omega'(a) := u_{r(a)} u_{c(a)}^{-1} a$. Since Ψ preserves the link grading, it follows that it intertwines Ω and Ω' , as desired. □

¹¹From a geometric viewpoint, \mathcal{S} indicates points where we add a π -rotation to the Lie group trivialization of the stabilized tangent bundle to \mathbb{S}^1 . Doing this an odd number of times on each component yields the null-cobordant trivialization. See [EES05, Remark 4.35].

Link automorphisms will appear in our discussion in two related ways. First, they naturally arise when considering the family of augmentations induced by an exact Lagrangian filling, as we will next describe in Section 3.3. Second, in Section 4.2 below, we describe a formula for the cobordism map over \mathbb{Z} associated to a saddle cobordism; our proof that the formula is correct is indirect and essentially reduces to arguing that there is only one possible candidate for the cobordism map that is actually a chain map over \mathbb{Z} . However, the existence of link automorphisms forces us to qualify this statement, since composing a chain map with a link automorphism produces another chain map. See Proposition 4.8 and Appendix A.

3.3. The geometric map induced by an exact Lagrangian cobordism. Suppose that Λ is a Legendrian link in $(\mathbb{R}^3, \xi_{\text{st}})$ and that $L \subset (\mathbb{R}^4, \lambda_{\text{st}})$ is a Lagrangian filling of Λ . Then L induces an augmentation of the Legendrian contact DGA $(\mathcal{A}_\Lambda, \partial)$. More precisely, the filling L equipped with a rank 1 local system induces an augmentation; put another way, the filling gives a family of augmentations and the additional choice of a local system picks out one of these. In the setting of $(\mathbb{R}^3, \xi_{\text{st}})$, the study of augmentations coming from fillings was initiated by Ekholm, Honda, and Kálmán [EHK16], who proved that an exact filling induces an augmentation over the group ring $\mathbb{Z}_2[H_1(L)]$ through a count of rigid holomorphic disks in the symplectization of \mathbb{R}^3 with boundary on L . Karlsson [Kar20] subsequently lifted \mathbb{Z}_2 to \mathbb{Z} by showing that the relevant moduli spaces of holomorphic disks can be coherently oriented. We summarize all of this work as follows.

Theorem 3.6 ([EHK16, Kar20]). *Suppose that L is an (oriented, embedded, exact) Lagrangian filling of the Legendrian link $\Lambda \subset (\mathbb{R}^3, \xi_{\text{st}})$ with Maslov number 0. Then L induces a DGA map*

$$\varepsilon_L : (\mathcal{A}_\Lambda, \partial) \rightarrow (\mathbb{Z}[H_1(L)], 0)$$

where $\mathbb{Z}[H_1(L)]$ lies entirely in grading 0. (The map ε_L is referred to as an augmentation.) Furthermore, if L and L' are Lagrangian fillings of Λ which are isotopic through exact Lagrangian fillings of Λ , then the corresponding augmentations ε_L and $\varepsilon_{L'}$ are DGA homotopic maps. \square

Note that an exact Lagrangian isotopy extends to an ambient Hamiltonian isotopy, e.g. by [Oh15, Section 3.6], especially [Oh15, Theorem 3.6.7], and see also [Pol01, Exercise 6.1.A]. Conversely, the image of an exact Lagrangian submanifold under a Hamiltonian diffeomorphism remains exact, and thus exact Lagrangian isotopies are equivalent to Hamiltonian isotopies. In fact, this also holds with compact support: [Oh15, Theorem 3.6.7] implies that a compactly supported exact Lagrangian isotopy extends to a compactly supported Hamiltonian isotopy.

Remark 3.7. For the definition of DGA homotopic maps, see e.g. [Kál05, EHK16, NRS⁺20]; we omit the definition here because for the Legendrian links that we consider in this paper, we can replace “DGA homotopic” by “the same”. All of our links Λ have rotation number 0 on each component, and all of the fillings that we construct are composed of minimum cobordisms and saddle cobordisms at Reeb chords with degree 0. It follows that each of these fillings has Maslov number 0. In addition, for all choices of Λ in this paper, all Reeb chords lie in nonnegative degree (in fact, in degree 0 or 1), and so \mathcal{A}_Λ is supported entirely in nonnegative degree. In this setting, two DGA maps $(\mathcal{A}_\Lambda, \partial) \rightarrow (\mathbb{Z}[H_1(L)], 0)$ are DGA homotopic if and only if they are equal. Thus if two fillings L, L' produce augmentations to $\mathbb{Z}[H_1(L)]$ that are distinct (under an isomorphism identifying $H_1(L)$ and $H_1(L')$), then we can use Theorem 3.6 to conclude that L, L' are not exact Lagrangian isotopic (or, equivalently in this setting, not Hamiltonian isotopic). \square

Remark 3.8. The augmentation ε_L depends on a choice of spin structure on the filling L , as explained in [Kar20]. If we change the spin structure by an element $\vartheta \in H^1(L; \mathbb{Z}_2)$, then we can define an isomorphism $\mathbb{Z}[H_1(L)] \rightarrow \mathbb{Z}[H_1(L)]$ by $x \mapsto (-1)^{\vartheta(x)}x$, and the augmentation

changes by composition with this isomorphism. This does not change the augmentation up to equivalence, in the sense of Definition 3.9 below. \square

It will be convenient for us to enlarge the coefficient ring $\mathbb{Z}[H_1(L)]$ to incorporate link automorphisms, as introduced in Section 3.2 above. Suppose that Λ is a Legendrian link with m components. Recall that given units u_1, \dots, u_m , we can define a link automorphism $\Omega : (\mathcal{A}_\Lambda, \partial) \rightarrow (\mathcal{A}_\Lambda, \partial)$. Any augmentation of $(\mathcal{A}_\Lambda, \partial)$ can be composed with this link automorphism to produce another augmentation, and so a single augmentation produces an $(m-1)$ -parameter family of augmentations. This family is parametrized by s_1, \dots, s_{m-1} , where we define $s_i = u_i/u_m$ for $i \leq m-1$. We restate this observation as follows.

Consider the ring $\mathbb{Z}[H_1(L)][s_1^{\pm 1}, \dots, s_{m-1}^{\pm 1}] \cong \mathbb{Z}[H_1(L) \oplus \mathbb{Z}^{m-1}]$. Then the augmentation $\varepsilon_L : (\mathcal{A}_\Lambda, \partial) \rightarrow (\mathbb{Z}[H_1(L)], 0)$ lifts to an augmentation

$$\tilde{\varepsilon}_L : (\mathcal{A}_\Lambda, \partial) \rightarrow (\mathbb{Z}[H_1(L) \oplus \mathbb{Z}^{m-1}], 0)$$

defined as follows: for any Reeb chord a of Λ ending on component $r(a)$ and beginning on component $c(a)$, we define $\tilde{\varepsilon}_L(a) := u_{r(a)}u_{c(a)}^{-1}\varepsilon_L(a)$, where $u_i = s_i$ for $i \leq m-1$ and $u_m = 1$.

The augmentation $\tilde{\varepsilon}_L$ to $\mathbb{Z}[H_1(L) \oplus \mathbb{Z}^{m-1}]$ incorporates both the geometry of the filling L and link automorphisms; henceforth we will view it as “the” augmentation coming from the filling L and will drop the tilde. We will also not need the distinction between generators of $H_1(L)$ and generators of \mathbb{Z}^{m-1} . It is then convenient to recast the augmentation ε_L in the following definition.

Definition 3.9. A k -system of augmentations of Λ is an algebra map

$$\varepsilon : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}[s_1^{\pm 1}, \dots, s_k^{\pm 1}]$$

such that $\varepsilon \circ \partial = 0$. By definition, two k -systems of augmentations

$$\varepsilon : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}[s_1^{\pm 1}, \dots, s_k^{\pm 1}], \quad \varepsilon' : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}[(s'_1)^{\pm 1}, \dots, (s'_k)^{\pm 1}]$$

are considered to be *equivalent* if there exists a \mathbb{Z} -algebra isomorphism

$$\psi : \mathbb{Z}[s_1^{\pm 1}, \dots, s_k^{\pm 1}] \rightarrow \mathbb{Z}[(s'_1)^{\pm 1}, \dots, (s'_k)^{\pm 1}]$$

such that $\varepsilon' = \psi \circ \varepsilon$. Note that the space of such isomorphisms is parametrized by $\mathbb{Z}_2^k \times \mathrm{GL}_k(\mathbb{Z})$. \square

Finally, we now recast Theorem 3.6 for our purposes in the following proposition; note that if L has genus g then $H_1(L) \oplus \mathbb{Z}^{m-1}$ has rank $2g + 2m - 2$.

Proposition 3.10. *Let Λ be an m -component Legendrian link. Let L be a connected, orientable exact Lagrangian filling of Λ of genus g and Maslov number 0. Then L gives rise to a $(2g + 2m - 2)$ -system of augmentations of Λ , and this system is well-defined, independent of choices, up to equivalence. Furthermore, if all Reeb chords of Λ have nonnegative degree, then isotopic fillings of Λ give rise to equivalent systems of augmentations. \square*

3.4. Signs and functoriality of the cobordism map. In order to establish our main results, such as Theorem 1.1, we will apply Proposition 3.10 to systems of augmentations that we will explicitly compute for particular fillings. For that, we will divide our fillings into elementary cobordism pieces, calculate the cobordism map for each elementary piece, and compose the resulting cobordism maps, using the fact that the cobordism map is functorial. This functoriality over \mathbb{Z} is established in the work of Karlsson [Kar20], and we summarize in this subsection the results from [Kar20] that we need.

Given an orientable exact Lagrangian cobordism L between Λ_+ and Λ_- , we choose a spin structure on L that restricts on each component of Λ_+ and Λ_- to the null-cobordant spin structure. Note that there are $|H_1(\widehat{L}; \mathbb{Z}_2)|$ such spin structures, where \widehat{L} is the closed surface obtained from L by gluing a disk to each boundary component, and any such spin structure

will do. Besides a spin structure on L , the other pieces of auxiliary data that Karlsson uses to define the cobordism maps are systems of capping operators for Λ_+ and Λ_- satisfying certain technical conditions. These capping operators are used by Karlsson to define the signs in the DGAs $(\mathcal{A}_{\Lambda_+}, \partial_+^{\text{NC}})$ and $(\mathcal{A}_{\Lambda_-}, \partial_-^{\text{NC}})$, where $\partial_{\pm}^{\text{NC}}$ are the differentials associated to the null-cobordant spin structures on Λ_{\pm} , as well as the signs in the cobordism map between the DGAs.

In our setting, for any Legendrian Λ with the Lie group spin structure, a suitable system of capping operators has been constructed in [EES05, Section 4.5], compare [Kar20, Remark 2.9]. These capping operators give precisely the signs for the DGA differential on Λ that we have presented combinatorially in Section 3.1 above and written as ∂^{comb} , see Remark 3.2. However, for the cobordism maps we need the signs from the null-cobordant rather than the Lie group spin structure. As explained in Remark 3.2, we can express this combinatorially by choosing a set \mathcal{S} of marked points on Λ with an odd number of marked points on each component, resulting in a differential $\partial^{\mathcal{S}}$ on \mathcal{A}_{Λ} such that we have an isomorphism

$$\phi^{\mathcal{S}} : (\mathcal{A}_{\Lambda}, \partial^{\mathcal{S}}) \xrightarrow{\cong} (\mathcal{A}_{\Lambda}, \partial^{\text{NC}}).$$

To return to the setting of a cobordism L between Λ_+ and Λ_- , Theorem 2.5 in [Kar20] gives a DGA map over \mathbb{Z} , $\Phi_L : (\mathcal{A}_{\Lambda_+}, \partial^{\text{NC}}) \rightarrow (\mathcal{A}_{\Lambda_-}, \partial^{\text{NC}})$. If we choose sets of marked points \mathcal{S}_{\pm} on Λ_{\pm} with an odd number of marked points on each component of Λ_{\pm} , then Φ_L induces a DGA map from $(\mathcal{A}_{\Lambda_+}, \partial^{\mathcal{S}_+})$ to $(\mathcal{A}_{\Lambda_-}, \partial^{\mathcal{S}_-})$. We also denote this map by Φ_L , and it satisfies that the following diagram commutes:

$$\begin{array}{ccc} (\mathcal{A}_{\Lambda_+}, \partial^{\mathcal{S}_+}) & \xrightarrow[\cong]{\phi^{\mathcal{S}_+}} & (\mathcal{A}_{\Lambda_+}, \partial^{\text{NC}}) \\ \Phi_L \downarrow & & \downarrow \Phi_L \\ (\mathcal{A}_{\Lambda_-}, \partial^{\mathcal{S}_-}) & \xrightarrow[\cong]{\phi^{\mathcal{S}_-}} & (\mathcal{A}_{\Lambda_-}, \partial^{\text{NC}}). \end{array}$$

Furthermore, the cobordism maps Φ_L constructed by Karlsson are functorial. To state this property, suppose that L_1 and L_2 are exact Lagrangian cobordisms that go from Λ_0 to Λ_1 and from Λ_1 to Λ_2 (from bottom to top), respectively. We can concatenate these to produce an exact cobordism $L_1 \# L_2$ from Λ_0 to Λ_2 . As before, equip L_1, L_2 with spin structures that restrict to the null-cobordant spin structures on their boundaries. Choices of capping operators on $\Lambda_0, \Lambda_1, \Lambda_2$ now produce DGA maps $\Phi_{L_1} : (\mathcal{A}_{\Lambda_1}, \partial^{\text{NC}}) \rightarrow (\mathcal{A}_{\Lambda_0}, \partial^{\text{NC}})$, $\Phi_{L_2} : (\mathcal{A}_{\Lambda_2}, \partial^{\text{NC}}) \rightarrow (\mathcal{A}_{\Lambda_1}, \partial^{\text{NC}})$, and $\Phi_{L_1 \# L_2} : (\mathcal{A}_{\Lambda_2}, \partial^{\text{NC}}) \rightarrow (\mathcal{A}_{\Lambda_0}, \partial^{\text{NC}})$, and [Kar20, Theorem 2.6] states that:

$$\Phi_{L_1} \circ \Phi_{L_2} = \Phi_{L_1 \# L_2}.$$

Let us choose collections of marked points $\mathcal{S}_0, \mathcal{S}_1, \mathcal{S}_2$ on $\Lambda_0, \Lambda_1, \Lambda_2$ such that each component has an odd number of marked points (as usual). Then, we can use the isomorphisms between $(\mathcal{A}_{\Lambda_i}, \partial^{\mathcal{S}_i})$ and $(\mathcal{A}_{\Lambda_i}, \partial^{\text{NC}})$ to produce DGA maps $\Phi_{L_1}, \Phi_{L_2}, \Phi_{L_1 \# L_2}$ between the DGAs $(\mathcal{A}_{\Lambda_i}, \partial^{\mathcal{S}_i})$ such that the following diagram commutes:

$$(3.1) \quad \begin{array}{ccc} & (\mathcal{A}_{\Lambda_2}, \partial^{\mathcal{S}_2}) \xrightarrow[\cong]{\phi^{\mathcal{S}_2}} (\mathcal{A}_{\Lambda_2}, \partial^{\text{NC}}) & \\ & \Phi_{L_2} \downarrow & \downarrow \Phi_{L_2} \\ \Phi_{L_1 \# L_2} \curvearrowleft & (\mathcal{A}_{\Lambda_1}, \partial^{\mathcal{S}_1}) \xrightarrow[\cong]{\phi^{\mathcal{S}_1}} (\mathcal{A}_{\Lambda_1}, \partial^{\text{NC}}) & \curvearrowright \Phi_{L_1 \# L_2} \\ & \Phi_{L_1} \downarrow & \downarrow \Phi_{L_1} \\ & (\mathcal{A}_{\Lambda_0}, \partial^{\mathcal{S}_0}) \xrightarrow[\cong]{\phi^{\mathcal{S}_0}} (\mathcal{A}_{\Lambda_0}, \partial^{\text{NC}}). & \end{array}$$

Note that all of the horizontal maps in this diagram are algebra maps over the relevant coefficient ring. Colloquially, they send each homology coefficient s_i to s_i , and not to $-s_i$.

This discussion above is summarized in the following result.

Proposition 3.11. *Given an exact Lagrangian cobordism L between Λ_+ and Λ_+ , and choices of marked points \mathcal{S}_\pm on Λ_\pm with an odd number on each component, we can write the cobordism map Φ_L as a DGA map from $(\mathcal{A}_{\Lambda_+}, \partial^{\mathcal{S}_+})$ to $(\mathcal{A}_{\Lambda_-}, \partial^{\mathcal{S}_-})$. If we have exact cobordisms L_1 from Λ_0 to Λ_1 and L_2 from Λ_1 to Λ_2 , and marked points $\mathcal{S}_0, \mathcal{S}_1, \mathcal{S}_2$ on $\Lambda_0, \Lambda_1, \Lambda_2$ with an odd number on each component, then the cobordism maps for L_1, L_2 , and their concatenation $L_1 \# L_2$ satisfy $\Phi_{L_1} \circ \Phi_{L_2} = \Phi_{L_1 \# L_2}$. \square*

3.5. System of augmentations for a decomposable filling. All the Lagrangian fillings that we consider in this paper are decomposable in the sense of [EHK16] (see Section 2.1). For a decomposable filling, one can explicitly construct the corresponding system of augmentations by composing the cobordism maps induced by each of the elementary cobordisms. These elementary cobordism maps are described in Sections 4.1 and 4.2 below. To combine them into the desired system of augmentations, we additionally need to keep track of base points and discuss how they produce the parameters in the system of augmentations. This is the content of the discussion that now follows.¹²

First, consider a general exact Lagrangian cobordism L between Legendrians Λ_+ and Λ_- , inducing a chain map Φ_L between the DGAs of Λ_+ and Λ_- . Recall from Section 3.1 that in the setting of the DGA of a Legendrian Λ , it is convenient to choose base points on Λ and use these points to keep track of the homology classes of the boundaries of the holomorphic disks that contribute to the differential. In a similar manner, we will keep track of homology classes contributing to Φ_L by placing arcs on L and counting intersections of holomorphic disks with these arcs.

To this end, suppose that we have a collection of oriented arcs and circles on L , such that all circles lie in the interior of L , the endpoints of all arcs lie on $\Lambda_+ \cup \Lambda_-$, and the arcs are transverse to $\Lambda_+ \cup \Lambda_-$ at their endpoints. Label these arcs $\gamma_1, \dots, \gamma_k$. Some subset $\{\gamma_{i_1}, \dots, \gamma_{i_p}\}$ has at least one endpoint on Λ_+ , and we view these endpoints as base points on Λ_+ ; similarly some subset $\{\gamma_{j_1}, \dots, \gamma_{j_q}\}$ has at least one endpoint on Λ_- , and we view these endpoints as base points on Λ_- . The chain map Φ_L between the DGAs of Λ_+ and Λ_- is defined by counting a finite collection of holomorphic disks with boundary on L and boundary punctures mapping to Reeb chords for Λ_+ and Λ_- ; we make the (generic) assumption that our curves γ_i intersect the boundaries of these disks transversely, and that no endpoint of an arc γ_i lies at the endpoint of a Reeb chord of Λ_+ or Λ_- .

In this setting, Φ_L is a map of algebras over the coefficient ring $\mathbb{Z}[s_1^{\pm 1}, \dots, s_k^{\pm 1}]$. More precisely, the DGA for Λ_+ equipped with the base points from $\gamma_{i_1}, \dots, \gamma_{i_p}$ has coefficient ring $\mathbb{Z}[s_{i_1}^{\pm 1}, \dots, s_{i_p}^{\pm 1}]$, and we can tensor this DGA over $\mathbb{Z}[s_{i_1}^{\pm 1}, \dots, s_{i_p}^{\pm 1}]$ with $\mathbb{Z}[s_1^{\pm 1}, \dots, s_k^{\pm 1}]$ to obtain a DGA over $\mathbb{Z}[s_1^{\pm 1}, \dots, s_k^{\pm 1}]$, which we write as $(\mathcal{A}_{\Lambda_+}, \partial_+)$. Similarly we can define the DGA $(\mathcal{A}_{\Lambda_-}, \partial_-)$ over $\mathbb{Z}[s_1^{\pm 1}, \dots, s_k^{\pm 1}]$. Then we can define the chain map $\Phi_L : \mathcal{A}_{\Lambda_+} \rightarrow \mathcal{A}_{\Lambda_-}$ as a map of $\mathbb{Z}[s_1^{\pm 1}, \dots, s_k^{\pm 1}]$ -algebras: each holomorphic disk Δ contributing to Φ_L is given the coefficient $s_1^{n_1(\Delta)} \dots s_k^{n_k(\Delta)} \in \mathbb{Z}[s_1^{\pm 1}, \dots, s_k^{\pm 1}]$, where $n_i(\Delta)$ counts the number of signed intersections of $\partial\Delta$ with the curve γ_i .

We now apply this discussion to describe how to concretely construct a system of augmentations for a Legendrian link Λ associated to a connected, decomposable exact Lagrangian filling L of Λ . Let m denote the number of components of Λ and g the genus of L . By assumption, L is a union of 0-handles (minimum cobordisms) and 1-handles (saddle cobordisms); let k denote the number of 0-handles, and note that it follows that there are $2g + m + k - 2 =: \ell$ 1-handles. We can cut off a small neighborhood of each minimum of L to produce a new

¹²We note that a similar treatment of base points on Lagrangian cobordisms and the induced DGA maps (over \mathbb{Z}_2) appears in [GSW20a, section 2].

cobordism L' whose top end is Λ and whose bottom end is a k -component unlink Λ_0 , such that L' is assembled out of just the 1-handles of L .

We can view L' through slices from top to bottom, so that it becomes a movie of embedded Legendrian links (except at finitely many times) starting with Λ , at the top, and ending with the k -component unlink Λ_0 , at the bottom. In the Lagrangian projection, each saddle move is then represented by replacing a (contractible) crossing by its 0-resolution. We can now add base points to this movie as follows. Place base points t_1, \dots, t_m on the m components of Λ . Each time we pass through a saddle, add two more base points labeled s_i and $-s_i^{-1}$. All base points persist to the bottom of the cobordism, Λ_0 . See Figure 16.




FIGURE 16. On the left, placing a pair of base points at the bottom of a saddle cobordism, representing opposite sides of an arc passing through the saddle point in the cobordism. On the right, dividing a decomposable filling of Λ into elementary pieces: from top to bottom, a sequence of saddle cobordisms ending at an unlink Λ_0 , and then filling in each unknot component.

On the Lagrangian cobordism L' , the base points t_1, \dots, t_m trace out arcs joining Λ to Λ_0 , while for each $i = 1, \dots, \ell$, the base points $s_i, -s_i^{-1}$ together trace out an arc joining Λ_0 to itself. We call these arcs τ_1, \dots, τ_m and $\sigma_1, \dots, \sigma_\ell$, respectively. Orient the τ_i arcs upwards, and orient the σ_i arcs so that in each slice the arc is oriented upwards at the point labeled s_i and downwards at the point labeled $-s_i^{-1}$. This places the decomposable cobordism between Λ_0 and Λ in the general picture described above of a cobordism equipped with oriented arcs.

Label the slices of L' from bottom to top by $\Lambda_0, \Lambda_1, \dots, \Lambda_\ell = \Lambda$, and divide L' into saddle cobordisms L_1, \dots, L_ℓ , where L_j is the piece of L' between Λ_{j-1} and Λ_j ; note that the saddle of L_j is associated to the arc $\sigma_{\ell+1-j}$. Each Λ_j is equipped with a collection of base points each labeled by either t_i or $\pm s_i^{\pm 1}$. For $j = 1, \dots, \ell$, let $(\mathcal{A}_{\Lambda_j}, \partial^{\text{comb}})$ denote the DGA of Λ_j over $\mathbb{Z}[t_1^{\pm 1}, \dots, t_m^{\pm 1}, s_1^{\pm 1}, \dots, s_\ell^{\pm 1}]$ with the differential ∂^{comb} defined combinatorially as in

Section 3.1 (note that some of the s_i parameters may not correspond to base points of Λ_j and thus may not appear in the definition of ∂^{comb}).

We next relate the DGAs $(\mathcal{A}_{\Lambda_j}, \partial^{\text{comb}})$ to the discussion from Section 3.4. To this end, for each $j = 0, \dots, \ell$, we identify a subset \mathcal{S}_j of the base points on Λ_j such that each component of Λ_j contains an odd number of points in \mathcal{S}_j ; we abbreviate this condition by calling such a subset *odd-cardinality*. We define \mathcal{S}_j by backwards induction on j . Let \mathcal{S}_ℓ be the collection of all of the base points t_1, \dots, t_m on $\Lambda_\ell = \Lambda$, and note that this is odd-cardinality. Given \mathcal{S}_j , each base point on Λ_j descends to a corresponding base point on Λ_{j-1} , and so we may view \mathcal{S}_j as a collection of base points on Λ_{j-1} . On Λ_{j-1} , we can add to \mathcal{S}_j one more base point, from the two new base points labeled by $\pm s_{\ell+1-j}^{\pm 1}$, such that the resulting collection \mathcal{S}_{j-1} is odd-cardinality: if Λ_{j-1} has one more component than Λ_j , then the choice of this extra base point is forced by the odd-cardinality condition, while if Λ_{j-1} has one fewer component than Λ_j , then we can choose either.

The choice of base points \mathcal{S}_j on Λ_j produces a differential $\partial^{\mathcal{S}_j}$ on \mathcal{A}_{Λ_j} as follows: first remove the $-$ signs at the front of any base points on Λ_j labeled by $-s_i^{-1}$, so that all base points are labeled by t_i or $s_i^{\pm 1}$; then negate any base point in \mathcal{S}_j , and let $\partial^{\mathcal{S}_j}$ be the resulting combinatorial differential as in Remark 3.2. Note that each t_i is negated in this process, while exactly one of s_i or s_i^{-1} is negated, depending on which of these base points is in \mathcal{S}_j . Thus we can define a \mathbb{Z} -algebra isomorphism

$$\psi : \mathbb{Z}[t_1^{\pm 1}, \dots, t_m^{\pm 1}, s_1^{\pm 1}, \dots, s_\ell^{\pm 1}] \rightarrow \mathbb{Z}[t_1^{\pm 1}, \dots, t_m^{\pm 1}, s_1^{\pm 1}, \dots, s_\ell^{\pm 1}]$$

by $\psi(t_i) = -t_i$ for all $i = 1, \dots, m$ and $\psi(s_i) = \pm s_i$ for each $i = 1, \dots, \ell$ (with the sign determined by whether s_i or s_i^{-1} is in \mathcal{S}_0), which extends to a map $\psi : \mathcal{A}_{\Lambda_j} \rightarrow \mathcal{A}_{\Lambda_j}$ by specifying $\psi(a) = a$ for all Reeb chords a . This map ψ now intertwines the differentials ∂^{comb} and $\partial^{\mathcal{S}_j}$:

$$\psi : (\mathcal{A}_{\Lambda_j}, \partial^{\text{comb}}) \xrightarrow{\cong} (\mathcal{A}_{\Lambda_j}, \partial^{\mathcal{S}_j}).$$

We can combine this with the isomorphism $\phi^{\mathcal{S}_j} : (\mathcal{A}_{\Lambda_j}, \partial^{\mathcal{S}_j}) \xrightarrow{\cong} (\mathcal{A}_{\Lambda_j}, \partial^{\text{NC}})$ from Section 3.4 to obtain an isomorphism $\phi^{\mathcal{S}_j} \circ \psi$ from $(\mathcal{A}_{\Lambda_j}, \partial^{\text{comb}})$ to the DGA $(\mathcal{A}_{\Lambda_j}, \partial^{\text{NC}})$ with the null-cobordant spin structure.

Recall from Section 3.4 that since each \mathcal{S}_j is odd-cardinality, each cobordism L_j induces a cobordism map $\Phi_{L_j} : (\mathcal{A}_{\Lambda_j}, \partial^{\mathcal{S}_j}) \rightarrow (\mathcal{A}_{\Lambda_{j-1}}, \partial^{\mathcal{S}_{j-1}})$. By combining this with the isomorphism ψ , we can view the cobordism map as a map $(\mathcal{A}_{\Lambda_j}, \partial^{\text{comb}}) \rightarrow (\mathcal{A}_{\Lambda_{j-1}}, \partial^{\text{comb}})$, which we also write as Φ_{L_j} , so that the following diagram commutes:

$$\begin{array}{ccc} (\mathcal{A}_{\Lambda_j}, \partial^{\text{comb}}) & \xrightarrow[\cong]{\psi} & (\mathcal{A}_{\Lambda_j}, \partial^{\mathcal{S}_j}) \\ \Phi_{L_j} \downarrow & & \downarrow \Phi_{L_j} \\ (\mathcal{A}_{\Lambda_{j-1}}, \partial^{\text{comb}}) & \xrightarrow[\cong]{\psi} & (\mathcal{A}_{\Lambda_{j-1}}, \partial^{\mathcal{S}_{j-1}}). \end{array}$$

Similarly, we can view the cobordism map $\Phi_{L'}$ as a DGA map $(\mathcal{A}_\Lambda, \partial^{\text{comb}}) \rightarrow (\mathcal{A}_{\Lambda_0}, \partial^{\text{comb}})$. By the functoriality property from Proposition 3.11, we have

$$\Phi_{L'} = \Phi_{L_1} \circ \dots \circ \Phi_{L_\ell} : (\mathcal{A}_\Lambda, \partial^{\text{comb}}) \rightarrow (\mathcal{A}_{\Lambda_0}, \partial^{\text{comb}}).$$

We obtain the filling L of Λ from the cobordism L' by filling in the k components of the unlink Λ_0 with disjoint Lagrangian disks. Each disk filling produces a unique augmentation, as we record in the following statement.

Proposition 3.12. *Let U denote the standard Legendrian unknot with a collection of base points with labels l_1, \dots, l_p (where typically each label is of the form $\pm s_i^{\pm 1}$ or $\pm t_i^{\pm 1}$). If $l_1 \cdots l_p = -1$ then the DGA $(\mathcal{A}_U, \partial_U)$ has a unique augmentation.*

Proof. Let a denote the Reeb chord of U . If l_1, \dots, l_q are the base points on one lobe of the figure eight in $\Pi_{xy}(U)$ and l_{q+1}, \dots, l_p are the base points on the other, then $\partial_U(a) = l_1 \cdots l_q + l_p^{-1} \cdots l_{q+1}^{-1}$. The condition for ε to be an augmentation is that $\varepsilon(\partial_U(a)) = 0$, in which case ε is uniquely determined since $\varepsilon(a) = 0$ for grading reasons. \square

Now let w_1, \dots, w_k denote the product of the labels of the base points on each of the k components of the unlink Λ_0 , and write R for the ring

$$R = (\mathbb{Z}[t_1^{\pm 1}, \dots, t_m^{\pm 1}, s_1^{\pm 1}, \dots, s_\ell^{\pm 1}]) / (w_1 = \cdots = w_k = -1).$$

Then the filling of Λ_0 by disks yields an augmentation

$$\varepsilon_0 : \mathcal{A}_{\Lambda_0} \rightarrow R.$$

Composing with $\Phi_{L'}$ now gives the augmentation of Λ induced by L :

$$\Phi_L = \varepsilon_0 \circ \Phi_{L'} : \mathcal{A}_\Lambda \rightarrow R.$$

We will already call this Φ_L the *combinatorial* system of augmentations of Λ induced by L , even though it will not become fully combinatorial until we present the combinatorial cobordism maps for isotopy cylinders and saddle cobordisms in Section 4. This is to temporarily distinguish Φ_L from the *geometric* system of augmentations of Λ from Proposition 3.10. In fact, the two systems agree up to equivalence, as we will show next.

3.6. The systems of augmentations agree. In this subsection, we prove that the combinatorial and geometric systems of augmentations of a decomposable filling L are equivalent. This result generalizes a result of Y. Pan from [Pan17b, section 3], which uses \mathbb{Z}_2 coefficients and treats the case where Λ has a single component. We use the same notation as in the previous subsection: Φ_L is a map from \mathcal{A}_Λ to R , where R is the ring $R = (\mathbb{Z}[t_1^{\pm 1}, \dots, t_m^{\pm 1}, s_1^{\pm 1}, \dots, s_\ell^{\pm 1}]) / (w_1 = \cdots = w_k = -1)$ with w_1, \dots, w_k being words associated to the k minima of L . The desired equivalence is shown in the following result, which will be proven momentarily:

Proposition 3.13. *Suppose that the filling L of Λ is connected. Then we have $R \cong \mathbb{Z}[\mathbb{Z}^{2g+2m-2}]$ and consequently Φ_L is a $(2g + 2m - 2)$ -system of augmentations of Λ . Furthermore, up to equivalence, Φ_L agrees with the geometric system of augmentations from Proposition 3.10.*

The crucial consequence of Proposition 3.13 is that since geometric systems of augmentations are invariant under Hamiltonian isotopy of the filling, the same is true of the combinatorial system of augmentations Φ_L . This is the fact that will allow us to distinguish fillings through a combinatorial calculation of their augmentations. Indeed, the following result is a direct consequence of Propositions 3.10 and 3.13:

Proposition 3.14. *Let L be a connected filling of Λ , and suppose that all Reeb chords of Λ have nonnegative degree. Then the combinatorial system of augmentations Φ_L of \mathcal{A}_Λ is invariant, up to equivalence, under exact Lagrangian isotopy of L .* \square

The argument for Proposition 3.13 above occupies the remainder of this section.

Proof of Proposition 3.13. By functoriality, Φ_L and the system of augmentations from Proposition 3.10 agree over \mathbb{Z} . What we need to do is keep track of the homology coefficients that appear in the definitions of the two families of augmentations, and show that the two agree up to equivalence. Thus, we reduce mod 2 and work with group rings over \mathbb{Z}_2 . In the course of tracking the homology coefficients, we will see that the abelian group generated multiplicatively by $t_1, \dots, t_m, s_1, \dots, s_\ell$ with relations $w_1 = \cdots = w_k = 1$ is isomorphic to a free abelian group with $2g + 2m - 2$ generators, whence it will follow that $R \cong \mathbb{Z}_2[\mathbb{Z}^{2g+2m-2}]$.

As in Section 3.5, let τ_i and σ_i denote the oriented arcs on L' corresponding to t_i and s_i . The map Φ_L counts intersections with τ_i and σ_i ; what we will show is that these counts keep track of homology classes in $H_1(L)$ along with link automorphisms. If a is a degree-0 Reeb chord of Λ , let $\mathcal{M}(a)$ denote the moduli space of (rigid) holomorphic disks Δ with boundary on L and a single positive boundary puncture mapping to a . We may assume that L is generic, so that none of the minima of L lies on the boundary of a holomorphic disk in any of the $\mathcal{M}(a)$. Recall that L' is obtained from L by removing a neighborhood of each minimum of L . By making these neighborhoods sufficiently small, we may assume that the boundary of each of the holomorphic disks $\Delta \in \mathcal{M}(a)$ lies entirely in L' and does not intersect the negative boundary Λ_0 of L' : that is, $\partial\Delta$ is an oriented arc on L' with endpoints at the endpoints of a . The cobordism map $\Phi_{L'}$ is then given as follows, for all degree 0 Reeb chords a of Λ :

$$\Phi_{L'}(a) = \sum_{\Delta \in \mathcal{M}(a)} w(\Delta) \in \mathbb{Z}_2[t_1^{\pm 1}, \dots, t_m^{\pm 1}, s_1^{\pm 1}, \dots, s_\ell^{\pm 1}],$$

where

$$(3.2) \quad w(\Delta) = \prod_{i=1}^m t_i^{\#(\partial\Delta \cap \tau_i)} \prod_{i=1}^{\ell} s_i^{\#(\partial\Delta \cap \sigma_i)}.$$

By the discussion preceding the proposition, the augmentation Φ_L is the composition of $\Phi_{L'}$ with the quotient map

$$\mathbb{Z}_2[t_1^{\pm 1}, \dots, t_m^{\pm 1}, s_1^{\pm 1}, \dots, s_\ell^{\pm 1}] \rightarrow \mathbb{Z}_2[t_1^{\pm 1}, \dots, t_m^{\pm 1}, s_1^{\pm 1}, \dots, s_\ell^{\pm 1}] / (w_1 = \dots = w_k = 1).$$

We want to compare Φ_L with the geometric setup from Section 3.3. Recall from Theorem 3.6 that L induces an augmentation $\varepsilon_L : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}_2[H_1(L)]$. This map agrees over \mathbb{Z}_2 with Φ_L but the group-ring coefficients are given by:

$$\varepsilon_L(a) = \sum_{\Delta \in \mathcal{M}(a)} \exp([\overline{\partial\Delta}]) \in \mathbb{Z}_2[H_1(L)].$$

The notation here is as follows. Choose a capping path γ_a for each Reeb chord a of Λ : a path in the connected surface L whose endpoints are the same as the endpoints of γ_a . For each disk $\Delta \in \mathcal{M}(a)$, close up the arc $\partial\Delta$ by adding the reverse of γ_a to give a closed loop $\overline{\partial\Delta} = (\partial\Delta) \cup (-\gamma_a)$. Then $\overline{\partial\Delta}$ represents a homology class in $H_1(L)$, and we denote this class in $\mathbb{Z}_2[H_1(L)]$ by $\exp([\overline{\partial\Delta}])$ (the exponential changes addition to multiplication).

We specify particular capping paths γ_a as follows. For $i = 0, \dots, \ell-1$, let $L_{>i} := L_{i+1} \cup \dots \cup L_\ell$ denote the portion of L above Λ_i , and $L_{>\ell} := \Lambda$. Note that $L_{>\ell}$ has m components while $L_{>0}$ has 1 component, and there are exactly $m-1$ values of i for which $L_{>(i-1)}$ has 1 fewer component than $L_{>i}$. For notational simplicity we will assume that these are the largest possible values: $i = \ell - m + 2, \dots, \ell$. (A similar argument holds in general.) In this case the first $m-1$ saddle moves from the top are all cobordisms that merge components. The arcs $\sigma_1, \dots, \sigma_{m-1}$ are the cores of these 1-handle attachments, and we write $\sigma_1^\vee, \dots, \sigma_{m-1}^\vee$ for the corresponding cocores. (More explicitly, begin at the i -th saddle, place one point on each strand of the crossing above this saddle, and trace this pair of points upwards through the cobordism to Λ to produce σ_i^\vee .) The paths $\sigma_1^\vee, \dots, \sigma_{m-1}^\vee$ join the m components of Λ to each other. For each Reeb chord a of Λ , we can now choose the capping path γ_a to lie on $\Lambda \cup \sigma_1^\vee \cup \dots \cup \sigma_{m-1}^\vee$ and to avoid the base points t_1, \dots, t_ℓ on Λ . By construction, among the arcs $\tau_1, \dots, \tau_m, \sigma_1, \dots, \sigma_\ell$, the only ones that γ_a intersects are some subset of $\sigma_1, \dots, \sigma_{m-1}$ determined by which components of Λ contain the endpoints of a .

Since the arcs $\sigma_1^\vee, \dots, \sigma_{m-1}^\vee$ form a tree connecting the components of Λ , we can find units $u_1, \dots, u_m \in \mathbb{Z}_2[s_1^{\pm 1}, \dots, s_{m-1}^{\pm 1}]$ such that for each $i = 1, \dots, m-1$, if σ_i^\vee ends on component $r(i)$ and begins on component $c(i)$ of Λ , then $s_i = u_{r(i)} u_{c(i)}^{-1}$; furthermore, (u_1, \dots, u_m) are well-defined once we specify $u_m = 1$, and the induced map $\mathbb{Z}_2[s_1^{\pm 1}, \dots, s_{m-1}^{\pm 1}] \rightarrow \mathbb{Z}_2[u_1^{\pm 1}, \dots, u_{m-1}^{\pm 1}]$

is an isomorphism. (Concretely, for any i there is a unique path from component m to component i that traverses $\pm\sigma_{j_1}^\vee, \dots, \pm\sigma_{j_k}^\vee$, where the \pm signs denote orientation, and then u_i is given by $s_{j_1}^{\pm 1} \cdots s_{j_k}^{\pm 1}$.) It now follows by the construction of the capping paths γ_a that if a is any Reeb chord of Λ and $r(a), c(a) \in \{1, \dots, m\}$ are the components of the ending and beginning points of a , then

$$\prod_{i=1}^{m-1} s_i^{\#(\gamma_a \cap \sigma_i)} = u_{r(a)} u_{c(a)}^{-1}.$$

Suppose for now that L has exactly one minimum; the general case will be considered afterward. Then Λ_0 is a single-component unknot U , and the product of the labels of the base points on Λ_0 is $t_1 \cdots t_m$ since the s_i base points cancel in pairs. Note that the abelian group generated by $t_1, \dots, t_m, s_1, \dots, s_\ell$ with a single relation $t_1 \cdots t_m = 1$ is free on $m + \ell - 1 = 2g + 2m - 2$ generators; Φ_L maps to the ring $\mathbb{Z}_2[t_1^{\pm 1}, \dots, t_m^{\pm 1}, s_1^{\pm 1}, \dots, s_\ell^{\pm 1}]/(t_1 \cdots t_m = 1) \cong \mathbb{Z}_2[\mathbb{Z}^{2g+2m-2}]$.

When L has one minimum, the relative homology $H_1(L, \Lambda)$ is generated by $\sigma_m, \dots, \sigma_\ell$ and $\tau_2 - \tau_1, \dots, \tau_m - \tau_1$. (Strictly speaking all of these arcs end on Λ_0 ; we extend these arcs by adding arcs in the disk filling Λ_0 , so that any endpoint on Λ_0 is replaced by an endpoint at the minimum of L .) Since $H_1(L)$ is dual to $H_1(L, \Lambda)$, we can compute the homology class $[\overline{\partial\Delta}]$ for $\Delta \in \mathcal{M}(a)$ by counting intersections with the generating set of $H_1(L, \Lambda)$. To be precise, we can identify $\mathbb{Z}_2[H_1(L)] \cong \mathbb{Z}_2[t_2^{\pm 1}, \dots, t_m^{\pm 1}, s_m^{\pm 1}, \dots, s_\ell^{\pm 1}]$, and under this isomorphism we have

$$\exp[\overline{\partial\Delta}] = \prod_{i=2}^m t_i^{\#(\overline{\partial\Delta} \cap \tau_i) - \#(\overline{\partial\Delta} \cap \tau_1)} \prod_{i=m}^{\ell} s_i^{\#(\overline{\partial\Delta} \cap \sigma_i)} = (t_2 \cdots t_m)^{-\#(\partial\Delta \cap \tau_1)} \prod_{i=2}^m t_i^{\#(\partial\Delta \cap \tau_i)} \prod_{i=m}^{\ell} s_i^{\#(\partial\Delta \cap \sigma_i)}.$$

We now compare this to the formula for $w(\Delta)$ in equation (3.2):

$$\begin{aligned} w(\Delta)|_{t_1=(t_2 \cdots t_m)^{-1}} &= \left(\prod_{i=1}^m t_i^{\#(\partial\Delta \cap \tau_i)} \right) \Big|_{t_1=(t_2 \cdots t_m)^{-1}} \prod_{i=1}^{m-1} s_i^{\#(\partial\Delta \cap \sigma_i)} \prod_{i=m}^{\ell} s_i^{\#(\partial\Delta \cap \sigma_i)} \\ &= (t_2 \cdots t_m)^{-\#(\partial\Delta \cap \tau_1)} \prod_{i=2}^m t_i^{\#(\partial\Delta \cap \tau_i)} \prod_{i=1}^{m-1} s_i^{\#(\gamma_a \cap \sigma_i)} \prod_{i=m}^{\ell} s_i^{\#(\partial\Delta \cap \sigma_i)} \\ &= u_{r(a)} u_{c(a)}^{-1} \exp[\overline{\partial\Delta}], \end{aligned}$$

where in the second equality we use the fact that $\partial\Delta$ and γ_a have the same endpoints and σ_i is a separating curve in L . Now, we extend

$$\varepsilon_L : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}_2[H_1(L)] \cong \mathbb{Z}_2[t_2^{\pm 1}, \dots, t_m^{\pm 1}, s_m^{\pm 1}, \dots, s_\ell^{\pm 1}]$$

by a link automorphism to

$$\tilde{\varepsilon}_L : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}_2[H_1(L) \oplus \mathbb{Z}^{m-1}] \cong \mathbb{Z}_2[t_2^{\pm 1}, \dots, t_m^{\pm 1}, s_m^{\pm 1}, \dots, s_\ell^{\pm 1}, u_1^{\pm 1}, \dots, u_{m-1}^{\pm 1}]$$

defined by $\tilde{\varepsilon}_L(a) := u_{r(a)} u_{c(a)}^{-1} \varepsilon_L(a)$, as in Section 3.3. Then, we have

$$\tilde{\varepsilon}_L(a) = \sum_{\Delta \in \mathcal{M}(a)} u_{r(a)} u_{c(a)}^{-1} \exp([\overline{\partial\Delta}]) = \sum_{\Delta \in \mathcal{M}(a)} w(\Delta)|_{t_1=(t_2 \cdots t_m)^{-1}}.$$

That is, the following diagram commutes:

$$\begin{array}{ccc} \mathcal{A}_\Lambda & \xrightarrow{\tilde{\varepsilon}_L} & \mathbb{Z}_2[t_2^{\pm 1}, \dots, t_m^{\pm 1}, s_m^{\pm 1}, \dots, s_\ell^{\pm 1}, u_1^{\pm 1}, \dots, u_{m-1}^{\pm 1}] \\ & \searrow \Phi_L & \downarrow \cong \\ & & \mathbb{Z}_2[t_1^{\pm 1}, \dots, t_m^{\pm 1}, s_1^{\pm 1}, \dots, s_\ell^{\pm 1}]/(t_1 \cdots t_m = 1). \end{array}$$

This shows that the combinatorial cobordism map Φ_L and the geometric cobordism map $\tilde{\varepsilon}_L$ agree up to isomorphism when L has one minimum.

Now suppose that L has $k > 1$ minima. We claim that we can reduce to the above case of a single minimum. The arcs σ_i have endpoints at the minima; since \tilde{L} is connected, there is a spanning tree of $k - 1$ arcs that connects all of the minima to each other. For notational simplicity, we assume that these arcs are $\sigma_{\ell-k+2}, \dots, \sigma_\ell$. Now imagine deforming L by homotopy equivalence by successively contracting each arc $\sigma_{\ell-k+2}, \dots, \sigma_\ell$ to a point. The result is a new surface \tilde{L} with a single minimum, which inherits the arcs τ_1, \dots, τ_m and $\sigma_i, i \leq \ell - k + 1$. The geometric cobordism map ε_L is defined homologically and does not change when we replace L by \tilde{L} .

We now examine what happens to the cobordism map Φ_L as we pass from L to \tilde{L} . Recall that Φ_L is an augmentation taking values in the ring

$$\mathbb{Z}_2[t_1^{\pm 1}, \dots, t_m^{\pm 1}, s_1^{\pm 1}, \dots, s_\ell^{\pm 1}]/(w_1 = \dots = w_k = 1),$$

where for $j = 1, \dots, k$, w_j is the word given by the product of the arcs having an endpoint at the j -th minimum (each endpoint contributes $t_i^{\pm 1}$ or $s_i^{\pm 1}$ depending on the orientation of the corresponding arc at the minimum). At the step where we contract σ_i , note that s_i appears in exactly two words w_{i_1} and w_{i_2} corresponding to the endpoints of σ_i . We use the relation for one of these words, $w_{i_1} = 1$, to solve for s_i , and substitute into $w_{i_2} = 1$; the result is exactly the relation corresponding to the new minimum given by contracting σ_i . Once we have contracted all of $\sigma_{\ell-k+2}, \dots, \sigma_\ell$, we are left with a single word w for the unique remaining minimum, and this process gives an isomorphism between the coefficient ring $\mathbb{Z}_2[t_1^{\pm 1}, \dots, t_m^{\pm 1}, s_1^{\pm 1}, \dots, s_\ell^{\pm 1}]/(w_1 = \dots = w_k = 1)$ for L and the ring $\mathbb{Z}_2[t_1^{\pm 1}, \dots, t_m^{\pm 1}, s_1^{\pm 1}, \dots, s_{\ell-k+1}^{\pm 1}]/(w = 1)$ for \tilde{L} . In particular, note that the abelian group generated by $t_1, \dots, t_m, s_1, \dots, s_\ell$ with relations $w_1 = \dots = w_k = 1$ is again free on $2g + 2m - 2$ generators, just as in the case where L has one minimum.




FIGURE 17. Sliding the arc $\partial\Delta$ to avoid intersections with σ_i , and then contracting σ_i .

Now in \tilde{L} , the boundaries $\partial\Delta$ of some holomorphic disks may pass through the minimum. To restore transversality, we perturb each $\partial\Delta$ as follows: at the step where we contract σ_i , we homotop $\partial\Delta$ near any intersection with σ_i so that it wraps around one of the endpoints of σ_i instead; see Figure 17. This removes any intersections of $\partial\Delta$ with σ_i , and it does not change the word $w(\partial\Delta)$ as given in (3.2) because of the relations $w_j = 1$. The end result is the surface \tilde{L} where all boundaries $\partial\Delta$ are disjoint from the minimum of \tilde{L} , and we have reduced to the case of a single minimum. This completes the proof. \square

Remark 3.15. The above proof shows that the augmentation/cobordism map $\Phi_L : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}_2[t_1^{\pm 1}, \dots, t_m^{\pm 1}, s_1^{\pm 1}, \dots, s_\ell^{\pm 1}]/(w_1 = \dots = w_k = 1)$ sends the product $t_1 \cdots t_m$ to 1, since the product $w_1 \cdots w_k$ is equal to $t_1 \cdots t_m$: each σ arc contributes endpoints that cancel, and each τ arc ends at exactly one of the minima.

We can lift this statement to \mathbb{Z} coefficients: if \mathcal{A}_Λ is the DGA of Λ with the Lie group spin structure, then $\Phi_L : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}[t_1^{\pm 1}, \dots, t_m^{\pm 1}, s_1^{\pm 1}, \dots, s_\ell^{\pm 1}]/(w_1 = \dots = w_k = 1)$ sends $t_1 \cdots t_m$

to $(-1)^m$. This follows from a result of Levenson [Lev17] that any augmentation of \mathcal{A}_Λ to a field (whether or not it comes from a filling) must send $t_1 \cdots t_m$ to $(-1)^m$, whence this must be true of Φ_L . \square

4. COBORDISM MAPS FOR ELEMENTARY COBORDISMS

Given a decomposable Lagrangian filling L of a Legendrian link Λ , we have described in Section 3 the general theory of how to build a system of augmentations for L . In order to apply this theory, we will use combinatorial formulas for cobordism maps corresponding to elementary cobordisms, which we can then compose to produce a formula for the cobordism map of an arbitrary decomposable filling. Of the three elementary cobordisms in Section 2.1, we have already discussed the DGA map for a minimum cobordism; see Proposition 3.12. In this section we present combinatorial formulas for the cobordism maps for the other two elementary cobordisms: isotopy cobordisms and saddle cobordisms.

The map for an isotopy cobordism (Section 4.1) is not new and dates back originally to work of Kálmán [Kál05]. The map for a saddle cobordism (Section 4.2) occupies the bulk of Section 4, with some technical details postponed to Appendix A. It builds on work of Ekholm–Honda–Kálmán [EHK16], but introduces two new features:

1. A combinatorial lift to integer coefficients \mathbb{Z} ,
2. A formula that (even) over \mathbb{Z}_2 works for some saddle cobordisms (where the combinatorial EHK map over \mathbb{Z}_2 does not).

In order to lift the saddle cobordism map to \mathbb{Z} , rather than directly constructing explicit orientations of the relevant moduli spaces, we use an ad hoc argument that allows us to deduce signs for a particularly simple saddle cobordism from the fact, due to work of Karlsson [Kar20], that the map must be a chain map over \mathbb{Z} . In fact we conclude a slightly weaker result: namely, we show that the cobordism map agrees with our combinatorial formula up to a link automorphism. Nevertheless, this additional choice of link automorphism will not affect our computations, and the statement we obtain is sufficient for the purposes of calculating augmentations for fillings. This is explained in Section 4.3.

4.1. The cobordism map for a Legendrian isotopy. In this subsection we review the cobordism map for an isotopy cobordism. Suppose that Λ_+ and Λ_- are Legendrian links related by a Legendrian isotopy. There is then a quasi-isomorphism between the DGAs $(\mathcal{A}_{\Lambda_+}, \partial)$ and $(\mathcal{A}_{\Lambda_-}, \partial)$, as first constructed by Chekanov [Che02] over \mathbb{Z}_2 and then lifted to \mathbb{Z} in [ENS02]. More precisely, these quasi-isomorphisms are DGA maps that are constructed for certain elementary Legendrian isotopies, to be described below. Any general Legendrian isotopy can be broken down into a sequence of elementary isotopies, and we compose the DGA maps for the elementary pieces to produce a DGA map for the isotopy.

This picture fits in a natural way with cobordism maps. Given a Legendrian isotopy between Λ_+ and Λ_- , let L denote the corresponding Lagrangian cobordism between Λ_+ and Λ_- . Then Ekholm–Honda–Kálmán [EHK16, section 6.3] show that over \mathbb{Z}_2 , the cobordism map $\Phi_L : (\mathcal{A}_{\Lambda_+}, \partial) \rightarrow (\mathcal{A}_{\Lambda_-}, \partial)$ agrees with the DGA map associated to the isotopy; note that by functoriality, it suffices to show this when L is the cobordism for an elementary isotopy. This result was subsequently upgraded to \mathbb{Z} coefficients by the combined work of Kálmán [Kál05], who showed that the map of DGAs over \mathbb{Z} associated to an isotopy (a path in the space of Legendrian links) is invariant under homotopy of the path; Ekholm–Kálmán [EK08], who showed that over \mathbb{Z}_2 , this DGA map gives the differential for the Legendrian contact DGA of the Legendrian surface given by the lift of L ; and Karlsson [Kar20, section 6], who showed that one can assign signs to the differential of this Legendrian surface to induce signs for the cobordism map Φ_L . For our purposes, we summarize this work as follows.

Proposition 4.1 ([Kál05, EK08, EHK16, Kar20]). *Suppose that Λ_+ and Λ_- are related by an elementary Legendrian isotopy, with corresponding Lagrangian cobordism L . Choose base points, a spin structure, and capping operators on Λ_+ ; these induce, via the isotopy, a corresponding choice of base points, spin structure, and capping operators on Λ_- . Then, the cobordism map $\Phi_L : (\mathcal{A}_{\Lambda_+}, \partial) \rightarrow (\mathcal{A}_{\Lambda_-}, \partial)$ is equal to the DGA map for the isotopy as constructed in [Che02, ENS02].*

By “elementary Legendrian isotopy”, we will mean one of the following three isotopies between Legendrian links with base points, all described in terms of their xy projections:

- Base point moves: fix the xy projection and move a base point across a crossing,
- Reidemeister III moves (triple point moves),
- Reidemeister II moves.

Any Legendrian isotopy can be decomposed into these elementary isotopies, along with planar isotopies of the xy projection in \mathbb{R}^2 .

In the remainder of this subsection, we review the combinatorial formulas from [Che02, ENS02] for the DGA maps for elementary isotopies. As usual, to compute the cobordism map for a general Legendrian isotopy, we can divide the isotopy into elementary isotopies and compose the resulting cobordism maps.

4.1.1. *Base point moves.* Suppose that Λ and Λ' are Legendrian links that are related by a base point move: outside of a neighborhood of a Reeb chord a , their xy projections agree, and inside this neighborhood, a base point moves across the crossing. See Figure 18. Then the DGA map for this move is $\Psi : (\mathcal{A}_\Lambda, \partial) \rightarrow (\mathcal{A}_{\Lambda'}, \partial')$ defined as follows: Ψ acts as the identity on all Reeb chords besides a and on all base point variables including s , and

$$\Psi(a) = sa \quad (\text{left diagram}) \qquad \Psi(a) = as^{-1} \quad (\text{right diagram}).$$

Note that Ψ is an isomorphism, and the DGA map for the reverse of one of these base point moves is Ψ^{-1} .




FIGURE 18. A base point move.

We observe that if we move a base point (or collection of base points) all the way around a component of Λ until it returns to where it started, the corresponding automorphism of $(\mathcal{A}_\Lambda, \partial)$ is the identity map. (This uses the fact that the variable s associated to the base point commutes with Reeb chord generators of \mathcal{A}_Λ ; in the fully noncommutative setting where s does not commute with Reeb chords, the automorphism is conjugation by s .) As a consequence, when calculating the cobordism map for an isotopy cobordism L , we do not need to specify an arc on L joining corresponding base points on the ends of L , since any two choices of such an arc will yield the same map.

4.1.2. *Reidemeister III moves.* Suppose that Λ and Λ' are related by a Reidemeister III move: see Figure 19. There are two types of Reidemeister III moves, III_a (left diagram) and III_b (right diagram); these are called “Move II” and “Move I” in [ENS02], respectively, and “L1a” and “L1b” in [EHK16]. There is a one-to-one correspondence between the Reeb chords of Λ and Λ' , with the correspondence between the three crossings involved in the move shown in Figure 19. Under this identification, \mathcal{A}_Λ and $\mathcal{A}_{\Lambda'}$ are identical.




FIGURE 19. The two types of Reidemeister III moves.

The DGA map for Reidemeister III_a , which we will actually not need in this paper, is simply the identity map on \mathcal{A}_Λ . To describe the DGA map for Reidemeister III_b , let $\sigma \in \{\pm 1\}$ denote the product of the orientation signs of the three quadrants of $\Pi_{xy}(\Lambda)$ indicated in Figure 19: this is $+1$ or -1 depending on whether an even or odd number of those quadrants are shaded. Then the DGA map $\Psi : (\mathcal{A}_\Lambda, \delta) \rightarrow (\mathcal{A}_{\Lambda'}, \delta')$ is defined to be the identity on all Reeb chords except for a_1 and on all base point variables, and

$$\Psi(a_1) = a_1 + \sigma a_3 a_2.$$

4.1.3. Reidemeister II moves.




FIGURE 20. A Reidemeister II move.

The DGA maps for a Reidemeister II move are more involved than for the other elementary isotopies. Suppose that Λ and Λ' are related by a Reidemeister II move, with $\Pi_{xy}(\Lambda')$ having two more crossings than $\Pi_{xy}(\Lambda)$, as shown in Figure 20.

Let a_1, \dots, a_r be the Reeb chords of Λ , and let b_1, b_2 denote the two new Reeb chords of Λ' . Write $(\mathcal{A}_\Lambda, \partial)$ and $(\mathcal{A}_{\Lambda'}, \partial')$ for the DGAs of Λ and Λ' . Let $|b_1| = i = |b_2| + 1$ in $\mathcal{A}_{\Lambda'}$, and construct the stabilization $(S(\mathcal{A}_\Lambda), \partial)$ by adding two generators e_1, e_2 with $|e_1| = i = |e_2| + 1$ to \mathcal{A}_Λ and extending the differential ∂ by $\partial(e_1) = e_2, \partial(e_2) = 0$. There is a chain isomorphism $\Psi : \mathcal{A}_{\Lambda'} \rightarrow S(\mathcal{A}_\Lambda)$ whose definition we recall below. We can then compose Ψ^{-1} with the inclusion map $i : \mathcal{A}_\Lambda \rightarrow S(\mathcal{A}_\Lambda)$ to get a chain map $\Psi^{-1} \circ i : \mathcal{A}_\Lambda \rightarrow \mathcal{A}_{\Lambda'}$. In the other direction, we can compose Ψ with the projection map $p : S(\mathcal{A}_\Lambda) \rightarrow \mathcal{A}_\Lambda$ sending each generator of \mathcal{A}_Λ to itself and sending e_1, e_2 to 0, to get a chain map $p \circ \Psi : \mathcal{A}_{\Lambda'} \rightarrow \mathcal{A}_\Lambda$:

$$\begin{array}{c}
 \Psi^{-1} \circ i \\
 \curvearrowright \\
 \mathcal{A}_\Lambda \xleftarrow{i} S(\mathcal{A}_\Lambda) \xleftarrow{\Psi} \mathcal{A}_{\Lambda'} \\
 \curvearrowleft \\
 p \circ \Psi
 \end{array}$$

Then $\Psi^{-1} \circ i$ and $p \circ \Psi$ are the cobordism maps for the cobordisms from Λ' to Λ and from Λ to Λ' , respectively, induced by the Reidemeister II isotopy.

We will need the precise definition of Ψ from [ENS02], and we recall it now. Let a_1, \dots, a_r be the Reeb chords of Λ , ordered in increasing height. Inductively construct a sequence of algebra isomorphisms $\Psi_1, \Psi_2, \dots, \Psi_r : \mathcal{A}_{\Lambda'} \rightarrow S(\mathcal{A}_\Lambda)$ as follows. By inspecting Figure 20, we see that there is a bigon for Λ' with $+$ corner at b_1 and $-$ corner at b_2 , and so we can write

$\partial'(b_1) = \sigma b_2 + v$ where $\sigma \in \{\pm 1\}$ and v counts disks with $+$ corner in the leftmost quadrant at b_1 . The map Ψ_1 (written as Φ_0 in [ENS02]) is defined by

$$\Psi_1(b_1) = e_1 \quad \Psi_1(b_2) = \sigma(e_2 - v) \quad \Psi_1(a_\ell) = a_\ell.$$

Given $\Psi_{\ell-1}$, we define $\Psi_\ell = g_\ell \circ \Psi_{\ell-1}$, where $g_\ell : S(\mathcal{A}_\Lambda) \rightarrow S(\mathcal{A}_\Lambda)$ is the identity on all generators except a_ℓ , and

$$g_\ell(a_\ell) = a_\ell + H(\partial a_\ell - \Psi_{\ell-1} \partial' a_\ell).$$

Here H is the map on $S(\mathcal{A}_\Lambda)$ (a module map, not an algebra map) defined by $H(w) = 0$ if w is any word that either does not contain e_1 or e_2 , or for which the leftmost e_i appearing in w is e_1 , and $H(w_1 e_2 w_2) = (-1)^{|w_1|+1} w_1 e_1 w_2$ if w_1 does not contain e_1 or e_2 . Finally, $\Psi = \Psi_r$.

Noting that for each ℓ , $\Psi(a_\ell) = \Psi_\ell(a_\ell) = g_\ell(a_\ell)$, we can restate the definition of Ψ more succinctly as follows:

$$(4.1) \quad \begin{aligned} \Psi(b_1) &= e_1 \\ \Psi(b_2) &= \sigma(e_2 - v) \\ \Psi(a_\ell) &= a_\ell - H(\Psi \partial' a_\ell). \end{aligned}$$

This definition looks circular since Ψ occurs on the right hand side of the definition of $\Psi(a_\ell)$, but in fact the height ordering and Stokes' Theorem imply that for any ℓ , $\partial' a_\ell$ involves only $b_1, b_2, a_1, \dots, a_{\ell-1}$ and not $a_{\ell+1}, \dots, a_r$, and so (4.1) can be used to recursively define $\Psi(a_\ell)$. Note that the height ordering does not appear explicitly in (4.1); however, the existence of the height filtration means that the recursive definition (4.1) terminates and thus produces a well-defined result.

Remark 4.2. It follows from the definition of Ψ that the chain map $p \circ \Psi : \mathcal{A}_\Lambda \rightarrow \mathcal{A}_\Lambda$ has the following simple form:

$$(p \circ \Psi)(b_1) = 0 \quad (p \circ \Psi)(b_2) = -\sigma v \quad (p \circ \Psi)(a_\ell) = a_\ell.$$

□

This concludes our description of the DGA maps associated to isotopy cobordisms.

4.2. The cobordism map for a saddle cobordism. We now address the cobordism map associated to a saddle cobordism. Let Λ_+ be a Legendrian link with a contractible Reeb chord a of degree 0; contractible chords of even degree can be similarly treated with suitable modification to the grading. In the xy projection, replacing the crossing a by its oriented resolution yields a Legendrian link Λ_- , and we write L_a for the saddle cobordism between Λ_- and Λ_+ .

Our goal in this subsection is to write down a combinatorial formula for the cobordism map $\Phi_{L_a} : (\mathcal{A}_{\Lambda_+}, \partial) \rightarrow (\mathcal{A}_{\Lambda_-}, \partial)$. In [EHK16], Ekholm–Honda–Kálmán describe such a formula for this map over \mathbb{Z}_2 , subject to the assumption that the Reeb chord a is what they call “simple”. Our goal here is to describe the EHK map over \mathbb{Z} and for what we call “proper chords”, which are a different (and apparently larger) class of contractible Reeb chords than simple chords. The proof that our map is indeed the geometric cobordism map Φ_{L_a} (stated as Proposition 4.8 below) is deferred to Appendix A.

Recall from Section 3.1 that the Legendrian contact differential ∂ for a Legendrian Λ counts immersed disks with a single $+$ corner, where “immersed disk” in our terminology includes the condition that all punctures are mapped to single quadrants (i.e., all corners are convex). We will now need to consider more general disks, which we call *immersed disks with concave corners*. These are immersed disks where each boundary puncture is again mapped to a crossing of the Lagrangian projection $\Pi_{xy}(\Lambda)$, but where we now allow a neighborhood of each boundary puncture to be mapped to either a single quadrant at the crossing (a convex corner) or the union of three quadrants (a concave corner). As with convex corners, we can

label each concave corner as positive or negative, depending on whether 2 of the 3 quadrants covered by the corner are positive or negative, respectively.

Definition 4.3. A contractible Reeb chord a of Λ_+ is *proper* if the following condition holds. For any immersed disk Δ , possibly with concave corners, such that:

- Δ has a positive convex corner at some Reeb chord besides a ,
- Δ has at least one positive convex corner at a ,
- all other convex corners of Δ are negative, and
- the only possible concave corners of Δ are positive concave corners at a ,

then it must be the case that the (closure of the) boundary of Δ in $\Pi_{xy}(\Lambda_+)$ passes through the crossing $\Pi_{xy}(a)$ only once. That is, any immersed disk Δ with the given properties must have no concave corners, Δ must have exactly one positive convex corner at a , and the boundary of Δ never passes through a except at that corner. \square

Remark 4.4. We remark that being proper and being simple in the sense of [EHK16] are not the same. We refer to [EHK16, Definition 6.16] for the index condition that defines the latter property. In our language, the condition for a Reeb chord a to be simple can be restated as follows: for any immersed disk Δ such that Δ has k positive convex corners at a , a positive convex corner at some Reeb chord besides a , and all other corners (including concave corners) being negative, it must be the case that Δ has at least k concave corners.¹³ The contractible Reeb chord a_9 in Figure 22 below is proper but not simple; the necessity of considering saddle moves at Reeb chords like a_9 in this paper is what motivated our definition of proper chords. We do not know if all simple contractible chords must be proper. \square

All of the Reeb chords that we use in this paper to perform saddle cobordisms are contractible and proper. This is a consequence of the following result.

Proposition 4.5. *If $\beta \in \text{Br}_N^+$ is an admissible braid and a is a crossing of β such that $\beta \setminus \{a\}$ contains a half-twist, then as a Reeb chord of the (-1) -closure $\Lambda(\beta)$, a is contractible and proper.*

Proof. Contractibility has already been shown in Proposition 2.8; we need to show properness. Suppose that Δ is a disk as in Definition 4.3. Because Δ has a positive convex corner at a , it must be “thin” in the sense that it lies in the neighborhood of the Legendrian unknot that contains the satellite $\Lambda(\beta)$. The presence of the half-twist, and the fact that Δ has no concave corners in the half-twist, prevents Δ from passing through the half-twist. This forces Δ to be embedded in the neighborhood of the unknot, and so its boundary only passes through a once. \square

We will next present a formula for the map for a saddle cobordism at a Reeb chord a when a is contractible and proper. As in [EHK16], the key is to consider immersed disks with two $+$ corners, one of which is at a . We break these into two types.

Let a_i be a Reeb chord of Λ_+ not equal to a . Define $\Delta_a^\rightarrow(a_i)$, respectively $\Delta_a^\leftarrow(a_i)$, to be the set of immersed disks for Λ_+ , such that:

- all corners are convex, and there are exactly two positive corners, one at a_i and one at a ;
- at the corner at a , the orientation of Λ points toward, respectively away from (for $\Delta_a^\leftarrow(a_i)$), the disk.

¹³In conversation with T. Ekhholm, it emerged that there is a typo in [EHK16, Definition 6.16]: the inequality $\text{ind}(u) \geq k$ in that definition should be $\text{ind}(u) \geq k + 1$. The revised inequality corresponds to our condition of having at least k concave corners.




FIGURE 21. A disk in $\Delta_a^-(a_i)$ (left) and a disk in $\Delta_a^+(a_i)$ (right). For the disk on the left, $w_1(\Delta) = a_3a_1$ and $w_2(\Delta) = a_2$. For the disk on the right, $w_1(\Delta) = 1$ and $w_2(\Delta) = a_4$.

See Figure 21. For any $\Delta \in \Delta_a^-(a_i) \cup \Delta_a^+(a_i)$, we can define three quantities. One is the sign $\text{sgn}(\Delta) \in \{\pm 1\}$, which is the product of the orientation signs over all corners of Δ , multiplied by the signs of any base points traversed by the boundary of the disk. The other are two words $w_1(\Delta), w_2(\Delta) \in \mathcal{A}_{\Lambda_+}$, defined as follows: $w_1(\Delta)$ is the product of the $-$ corners and base points that we encounter as we traverse the boundary of Δ counterclockwise from a_i to a , and $w_2(\Delta)$ is the analogous product as we traverse the boundary counterclockwise from a to a_i . See Figure 21 for an example.

Definition 4.6. The combinatorial cobordism map, denoted $\Phi_{L_a}^{\text{comb}} : \mathcal{A}_{\Lambda_+} \rightarrow \mathcal{A}_{\Lambda_-}$, is the composition of three algebra maps:

$$\Phi_{L_a}^{\text{comb}} := \Phi^{\leftarrow} \circ \Phi^{\rightarrow} \circ \Phi^0,$$

where $\Phi_0 : \mathcal{A}_{\Lambda_+} \rightarrow \mathcal{A}_{\Lambda_-}$ is defined by $\Phi_0(a) = s$ and $\Phi_0(a_i) = a_i$ for any Reeb chord a_i besides a , and $\Phi^{\rightarrow}, \Phi^{\leftarrow} : \mathcal{A}_{\Lambda_-} \rightarrow \mathcal{A}_{\Lambda_-}$ are defined as follows. Let a_i be a generator of \mathcal{A}_{Λ_-} , that is, a Reeb chord of Λ_- , which is then also a Reeb chord of Λ_+ . Then,

$$\begin{aligned} \Phi^{\rightarrow}(a_i) &:= a_i + \sum_{\Delta \in \Delta_a^-(a_i)} (-1)^{|w_1(\Delta)|} \text{sgn}(\Delta) \Phi^{\rightarrow}(w_1(\Delta)) s^{-1} w_2(\Delta) \\ \Phi^{\leftarrow}(a_i) &:= a_i + \sum_{\Delta \in \Delta_a^+(a_i)} (-1)^{|w_1(\Delta)|} \text{sgn}(\Delta) \Phi^{\leftarrow}(w_1(\Delta)) s^{-1} w_2(\Delta). \end{aligned}$$

□

Remark 4.7. As with the definition of the Reidemeister II cobordism map, equation (4.1) in Section 4.1, these definitions may appear circular but can be used to recursively define Φ^{\rightarrow} and Φ^{\leftarrow} . The reason is that if we order the Reeb chords a_1, \dots, a_r in increasing order of height, then all disks with positive punctures at a and a_j can only have negative corners at a_1, \dots, a_{j-1} and not at a_{j+1}, \dots, a_r : in particular, if $\Delta \in \Delta_a^-(a_j) \cup \Delta_a^+(a_j)$ then $w_1(\Delta)$ only involves a_1, \dots, a_{j-1} . □

The key result in this subsection is the relation between $\Phi_{L_a}^{\text{comb}}$, as defined above, and Φ_{L_a} . This is the content of the following result:

Proposition 4.8. *If $a \in \mathcal{A}_{\Lambda_+}$ is a proper contractible Reeb chord, then the cobordism map $\Phi_{L_a} : \mathcal{A}_{\Lambda_+} \rightarrow \mathcal{A}_{\Lambda_-}$ is equal to the combinatorial map $\Phi_{L_a}^{\text{comb}}$, up to a link automorphism of Λ_- . That is, there is a link automorphism $\Omega : \mathcal{A}_{\Lambda_-} \rightarrow \mathcal{A}_{\Lambda_-}$ such that $\Phi_{L_a} = \Omega \circ \Phi_{L_a}^{\text{comb}}$.*

Proposition 4.8 is proved in Appendix A below. Let us illustrate how to compute $\Phi_{L_a}^{\text{comb}}$ in an explicit example, which will also appear as part of our later computations.

Example 4.9. Consider the configuration shown in Figure 22, this appears as part of our calculations for the \tilde{D}_4 -Legendrian in Section 6.2. The first step in that calculation is a saddle




FIGURE 22. Calculating the cobordism map for a saddle cobordism at a_9 . Left, the Legendrian link at the top of the cobordism; right, an immersed disk showing that a_9 is not simple.

move at a_9 , and we calculate the corresponding map Φ^{\leftarrow} here. By inspection we see that $\Delta_{a_9}^{\leftarrow}(a_i) = \emptyset$ for $10 \leq i \leq 13$, while $\Delta_{a_9}^{\leftarrow}(a_{14})$ and $\Delta_{a_9}^{\leftarrow}(a_{15})$ each contain one disk apiece, with negative corners at a_{13}, a_{11} and a_{10}, a_{13} respectively. For a_{16} , $\Delta_{a_9}^{\leftarrow}(a_{16})$ contains three disks, one with no negative corners, one with negative corners at a_{10}, a_{14} , and one with negative corners at a_{15}, a_{11} . It follows from this that $\Phi^{\leftarrow}(a_i) = a_i$ for $10 \leq i \leq 13$ and

$$\begin{aligned} \Phi^{\leftarrow}(a_{14}) &= a_{14} - \Phi^{\leftarrow}(a_{13}a_{11})s^{-1} = a_{14} - a_{13}a_{11}s^{-1} \\ \Phi^{\leftarrow}(a_{15}) &= a_{15} - \Phi^{\leftarrow}(1)s^{-1}a_{10}a_{13} = a_{15} - s^{-1}a_{10}a_{13} \\ \Phi^{\leftarrow}(a_{16}) &= a_{16} - s^{-1} - s^{-1}a_{10}a_{14} - \Phi^{\leftarrow}(a_{15}a_{11})s^{-1} = \\ &= a_{16} - s^{-1} - s^{-1}a_{10}a_{14} - (a_{15} - s^{-1}a_{10}a_{13})a_{11}s^{-1}. \end{aligned}$$

For the complete Legendrian that we study in Section 6.2, an inspection of Figure 28 shows that $\Delta_{a_9}^{\rightarrow}(a_i) = \emptyset$ and thus $\Phi^{\rightarrow}(a_i) = a_i$ for $10 \leq i \leq 16$. It follows that the map $\Phi_{L_{a_9}}^{\text{comb}}$ sends a_9 to s and agrees with Φ^{\leftarrow} for a_i , $10 \leq i \leq 16$.

We note that in this example, a_9 is contractible and proper but not simple, and thus even over \mathbb{Z}_2 we cannot directly apply the combinatorial formula from [EHK16]. The reason a_9 is not simple (cf. Remark 4.4) is the disk shown in Figure 22, which has 2 positive corners at a_9 , 1 positive corner at a_{16} , and a single concave corner at a_{13} . \square

Remark 4.10. If a is not just proper but also simple, then our definition of $\Phi_{L_a}^{\text{comb}}$ can be stated in an easier way, to match [EHK16]. In this case, write $\Delta_a(a_i) = \Delta_a^{\rightarrow}(a_i) \cup \Delta_a^{\leftarrow}(a_i)$. If Δ is any disk in $\Delta_a(a_i)$ and a_j is a negative corner of Δ , then it must be the case that $\Delta_a(a_j) = \emptyset$; otherwise the union of Δ and a disk $\Delta' \in \Delta_a(a_j)$ is an immersed disk with concave corner at a_j and two positive (convex) corners at a , violating the simplicity condition. Then we can drop the Φ^{\rightarrow} and Φ^{\leftarrow} in $\Phi^{\rightarrow}(w_1(\Delta))$ and $\Phi^{\leftarrow}(w_1(\Delta))$, and conclude directly that for all a_i ,

$$\Phi_{L_a}^{\text{comb}}(a_i) = a_i + \sum_{\Delta \in \Delta_a(a_i)} \text{sgn}(\Delta)w_1(\Delta)s^{-1}w_2(\Delta).$$

If we set $s = 1$ and reduce mod 2, this recovers the formula for the cobordism map from [EHK16, Proposition 6.18]. \square

4.3. Assembling elementary cobordism maps. Having described the cobordism maps for elementary cobordisms, we can calculate the map associated to any decomposable cobordism by composing the maps for its elementary pieces, and indeed this is what we do in Sections 5 and 6 below. There is a possible difficulty with this approach: we have only calculated the saddle cobordism map up to a link automorphism (see Proposition 4.8). However, for a filling, the extra flexibility provided by the basepoint parameters gets rid of this problem, as we explain in this subsection.

Let L be a connected decomposable genus- g filling of an m -component Legendrian link Λ . As in Section 3.5, we decorate L with arcs corresponding to base points $t_1, \dots, t_m, s_1, \dots, s_\ell$. Divide L into elementary cobordisms L_1, \dots, L_k , where:

- L_j is a cobordism between Legendrians Λ_{j-1} and Λ_j , with $\Lambda_0 = \emptyset$ and $\Lambda_k = \Lambda$;
- L_1 is a disjoint union of minimum cobordisms;
- for $j = 2, \dots, k$, L_j is either an isotopy cobordism or a saddle cobordism.

Note that this decomposition differs slightly from our simplified setup in Section 3.5, where we suppressed isotopy cobordisms.

As in Section 3.5, let R be the ring

$$R := (\mathbb{Z}[t_1^{\pm 1}, \dots, t_m^{\pm 1}, s_1^{\pm 1}, \dots, s_\ell^{\pm 1}]) / (w_1 = \dots = w_k = -1) \cong \mathbb{Z}[H_1(L) \oplus \mathbb{Z}^{m-1}],$$

where w_1, \dots, w_k are words coming from the minima of L . For each j , let $(\mathcal{A}_{\Lambda_j}, \partial^{\text{comb}})$ denote the DGA for Λ_j over R , with $\mathcal{A}_\Lambda = \mathcal{A}_{\Lambda_k}$. Then each elementary cobordism gives a map $\Phi_{L_j} : (\mathcal{A}_{\Lambda_j}, \partial^{\text{comb}}) \rightarrow (\mathcal{A}_{\Lambda_{j-1}}, \partial^{\text{comb}})$, and their composition is a $(2g + 2m - 2)$ -system of augmentations for Λ :

$$\Phi_L = \Phi_{L_1} \circ \dots \circ \Phi_{L_k} : (\mathcal{A}_\Lambda, \partial^{\text{comb}}) \rightarrow (R, 0).$$

Now suppose that for $j = 1, \dots, k$, $\Omega_j : \mathcal{A}_{\Lambda_{j-1}} \rightarrow \mathcal{A}_{\Lambda_j}$ is a link automorphism of Λ_{j-1} , and define $\tilde{\Phi}_{L_j} = \Omega_j \circ \Phi_{L_j}$.

Proposition 4.11. *The maps $\Phi_L = \Phi_{L_1} \circ \dots \circ \Phi_{L_k}$ and $\tilde{\Phi}_L = \tilde{\Phi}_{L_1} \circ \dots \circ \tilde{\Phi}_{L_k}$ are equivalent systems of augmentations of Λ .*

Proof. We prove by induction that for $j = 1, \dots, k$, $\Phi_{L_1} \circ \dots \circ \Phi_{L_j}$ and $\tilde{\Phi}_{L_1} \circ \dots \circ \tilde{\Phi}_{L_j}$ are equivalent as maps $(\mathcal{A}_{\Lambda_j}, \partial^{\text{comb}}) \rightarrow (R, 0)$. The base case $j = 1$ is true since Φ_{L_1} , and thus $\tilde{\Phi}_{L_1}$, are both the zero map on Reeb chords of Λ_1 .

For the induction step, assume that $\Phi_{L_1} \circ \dots \circ \Phi_{L_j}$ and $\tilde{\Phi}_{L_1} \circ \dots \circ \tilde{\Phi}_{L_j}$ are equivalent, so that there is an automorphism ψ_j of R such that $\tilde{\Phi}_{L_1} \circ \dots \circ \tilde{\Phi}_{L_j} = \psi_j \circ (\Phi_{L_1} \circ \dots \circ \Phi_{L_j})$. Since the map $\Phi_{L_1} \circ \dots \circ \Phi_{L_j}$ agrees with the geometric system of augmentations for $L_1 \cup \dots \cup L_j$ by Proposition 3.13, and the geometric system incorporates link automorphisms of Λ_j , the link automorphism Ω_j of Λ_j induces an automorphism ω_j of R such that

$$(\Phi_{L_1} \circ \dots \circ \Phi_{L_j}) \circ \Omega_j = \omega_j \circ (\Phi_{L_1} \circ \dots \circ \Phi_{L_j}).$$

(Note that Proposition 3.13 assumes that $L_1 \cup \dots \cup L_j$ is connected; however, the argument here extends to the disconnected case as well, since the system of augmentations of a disconnected filling annihilates any Reeb chord with endpoints on different components.)

We conclude that the following diagram commutes:

$$\begin{array}{ccccc}
& & \mathcal{A}_{\Lambda_{j+1}} & & \\
& \swarrow \Phi_{L_{j+1}} & & \searrow \tilde{\Phi}_{L_{j+1}} & \\
& \mathcal{A}_{\Lambda_j} & \xrightarrow{\Omega_j} & \mathcal{A}_{\Lambda_j} & \\
& \swarrow \Phi_{L_1} \circ \dots \circ \Phi_{L_j} & \xrightarrow{\cong} & \mathcal{A}_{\Lambda_j} & \searrow \tilde{\Phi}_{L_1} \circ \dots \circ \tilde{\Phi}_{L_j} \\
R & \xrightarrow{\omega_j} & R & \xrightarrow{\psi_j} & R.
\end{array}$$

It follows that $\Phi_{L_1} \circ \dots \circ \Phi_{L_j} \circ \Phi_{L_{j+1}}$ and $\tilde{\Phi}_{L_1} \circ \dots \circ \tilde{\Phi}_{L_j} \circ \tilde{\Phi}_{L_{j+1}}$ are equivalent since one is the composition of the other with $\psi_j \circ \omega_j$, and this completes the induction. \square

By Proposition 4.11, when we build systems of augmentations for fillings by composing elementary cobordism maps, we can replace any elementary cobordism map by its composition with a link automorphism. In particular, Proposition 4.8 implies that we can use the combinatorial saddle map Φ^{comb} as the cobordism map for a saddle cobordism, and this is what we will do in subsequent sections.

5. LEGENDRIAN CONTACT DGA AND COBORDISM MAPS FOR (-1) -CLOSURES

In this section we present an algebraically amenable description of the Legendrian contact DGA for the (-1) -closure of an admissible braid, and detail the effect of the ϑ -loops and saddle cobordisms on the DGA.

5.1. The DGA of the (-1) -closure of an admissible braid. Let $\sigma_{k_1} \cdots \sigma_{k_r} \in \text{Br}_N^+$ be an admissible positive braid. Henceforth we will write $\Lambda(\sigma_{k_1} \cdots \sigma_{k_r})$ for the (-1) -closure of this braid in the sense of Definition 2.3.¹⁴ We decorate the xy projection of $\Lambda(\sigma_{k_1} \cdots \sigma_{k_r})$ as follows; see Figure 23. Place a column of base points on the n strands of the braid between braid crossings, as well as on either end of the braid, and label these base points $t_{\ell,i}$, $1 \leq i \leq n$, $0 \leq \ell \leq r$. (In practice we may only need some small subset of these base points; in that case we formally set $t_{\ell,i} = 1$ for all of the other base points and then remove them.) The Reeb chords for $\Lambda(\sigma_{k_1} \cdots \sigma_{k_r})$ consist of:

- a_1, \dots, a_r , of degree 0, corresponding to the crossings of the braid, and labeled in the obvious way;
- c_{ij} , $1 \leq i, j \leq n$, of degree 1, corresponding to the Reeb chord of the standard Legendrian unknot U .

Recall from Section 3.1 that in order to calculate degrees of Reeb chords, we need to choose a base point on each component of the link; any subset of the $t_{\ell,i}$ will do and produces the degrees given above.




FIGURE 23. The Lagrangian projection of the Legendrian link $\Lambda(\sigma_{k_1} \cdots \sigma_{k_r})$, with crossings and base points labeled. The braid itself is in the blue box. Arrows represent the orientation of the link.

The differential on the Legendrian contact DGA of $\Lambda(\sigma_{k_1} \cdots \sigma_{k_r})$ can be expressed in a compact way using the path matrices of Kálmán [Kál06].¹⁵ For $k = 1, \dots, n - 1$, define an $n \times n$ matrix $P_k(a)$ (as a function of an input a) as follows:

$$(P_k(a))_{ij} = \begin{cases} 1 & i = j \text{ and } i \neq k, k + 1 \\ 1 & (i, j) = (k, k + 1) \text{ or } (k + 1, k) \\ a & i = j = k + 1 \\ 0 & \text{otherwise;} \end{cases}$$

¹⁴Note that this differs from the notation $\Lambda(\beta)$ in Section 2.2, but by Proposition 2.6, the two notations represent links that are Legendrian isotopic.

¹⁵Note that we number our braid strands in increasing order from bottom to top, while Kálmán numbers braid strands from top to bottom. We also incorporate base points while Kálmán does not.

that is, $P_k(a)$ is the identity matrix except for the 2×2 submatrix given by rows and columns k and $k+1$, which is $\begin{pmatrix} 0 & 1 \\ 1 & a \end{pmatrix}$. (These are the path matrices considered in [Kál06], but note that we number our braid strands in increasing order from bottom to top, while Kálmán numbers braid strands from top to bottom.) Also define $\mathbf{t}_\ell = (t_{\ell,1}, \dots, t_{\ell,n})$ and write $D(\mathbf{t}_\ell)$ for the diagonal $n \times n$ matrix with $t_{\ell,1}, \dots, t_{\ell,n}$ along the diagonal.

Definition 5.1. Let $\beta = \sigma_{k_1} \cdots \sigma_{k_r}$ be an n -stranded braid decorated with base points, with crossings and base points labeled as in Figure 23. The *path matrix* of β is the $n \times n$ matrix

$$P_\beta = D(\mathbf{t}_0)P_{k_1}(a_1)D(\mathbf{t}_1)P_{k_2}(a_2)D(\mathbf{t}_2) \cdots P_{k_r}(a_r)D(\mathbf{t}_r).$$

Colloquially, the (i, j) entry of the path matrix $P(\beta)$ counts paths beginning at the left of β on strand i , ending at the right on strand j , and at each crossing the path encounters, either passing straight through the crossing, or turning a corner if the path changes direction from southeast to northeast at the corner. Each path produces a word by reading the base points traversed and corners turned in order, and the (i, j) entry of $P(\beta)$ is the sum of these words.

Proposition 5.2. *The differential on the DGA $(\mathcal{A}_{\Lambda(\sigma_{k_1} \cdots \sigma_{k_r})}, \partial)$ for $\Lambda(\sigma_{k_1} \cdots \sigma_{k_r})$ is given as follows: $\partial(a_\ell) = 0$, and if we assemble the c_{ij} into an $n \times n$ matrix $C = (c_{ij})$ and write $\mathbf{1}$ for the $n \times n$ identity matrix, then:*

$$\partial(C) = \mathbf{1} + P_\beta.$$

Proof. Each degree-0 generator a_ℓ has vanishing differential for degree reasons. For c_{ij} , there are two possible types of immersed disks (all of which are in fact embedded) with $+$ corner at c_{ij} , depending on which $+$ quadrant at c_{ij} is covered by the disk. There is an embedded disk with $+$ puncture at the right quadrant of c_{ij} and no $-$ puncture if $i = j$, and otherwise there is no immersed disk with $+$ puncture at this right quadrant. This produces the $\mathbf{1}$ term in the formula. For embedded disks with $+$ puncture at the left quadrant of c_{ij} , we need to keep track of ways that the boundary of this disk can enter the braid from the left on strand i and exit the braid to the right on strand j , with possible convex corners at some crossings a_ℓ . The contribution of these disks to $\partial(c_{ij})$ is precisely the (i, j) entry of the path matrix P_β . \square

5.2. ϑ -monodromy action on the DGA. Consider a Legendrian link $\Lambda = \Lambda(\beta, k; \gamma) \subset (\mathbb{R}^3, \xi_{\text{st}})$ and its ϑ -loop, as defined in Section 2.4. Here we compute the morphism

$$\mathcal{A}(\vartheta) : \mathcal{A}_\Lambda \rightarrow \mathcal{A}_\Lambda$$

induced by this Legendrian isotopy, which we call the ϑ -monodromy or purple box monodromy. To be precise, any Legendrian isotopy between Legendrian links induces a chain isomorphism between the (suitably stabilized) DGAs of the links, as described in [Che02, ENS02]. In the case of the isotopy given by the ϑ -loop, which consists entirely of Reidemeister III moves, it is not necessary to stabilize the DGAs, and as a result we obtain the aforementioned chain isomorphism $\mathcal{A}(\vartheta)$, which we now compute explicitly.

The Lagrangian projection of $\Lambda(\beta, k; \gamma)$ is given in the right diagram in Figure 13. Let n denote the braid index of γ , and let N be the number of braid strands in $\Lambda(\beta, k; \gamma)$, so that the braid index of β is $N - n + 1$. As in Section 5.1, the Reeb chords of $\Lambda(\beta, i; \gamma)$, which generate the Legendrian contact DGA \mathcal{A}_Λ , come in two types: the degree 1 chords c_{ij} , $i, j \in [1, N]$, and the degree 0 chords in the braiding region. We can divide these Reeb chords into two types in another way. Call the sublink of $\Lambda(\beta, i; \gamma)$ corresponding to the i -th strand of β (and containing the purple box γ) the *satellite sublink*; this is depicted in purple in Figure 13. We call crossings of $\Lambda(\beta, i; \gamma)$ *satellite crossings* and *non-satellite crossings* depending on whether or not they involve the satellite sublink. Note that the satellite crossings of degree 1 are precisely c_{ij} with $i, j \in \{k, \dots, k+n-1\}$, while the satellite crossings of degree 0 come in groups of n , with each group coming from a single crossing of β .

We allow for the placement of arbitrarily many base points on $\Lambda(\beta, k; \gamma)$, subject to the restriction that any base points lying on the satellite sublink actually lie in the purple box for γ . (In practice, there will be one base point per strand of $\Lambda(\beta, k; \gamma)$, and the base points in the purple box will lie on its right edge.) Let P_γ denote the $n \times n$ path matrix for γ with its base points. Extend this to an $N \times N$ matrix \tilde{P}_γ by

$$\tilde{P}_\gamma = \begin{pmatrix} \mathbf{1} & 0 & 0 \\ 0 & P_\gamma & 0 \\ 0 & 0 & \mathbf{1} \end{pmatrix}$$

where the central matrix P_γ corresponds to rows and columns $k, \dots, k+n-1$.

Proposition 5.3. *The purple-box monodromy map $\mathcal{A}(\vartheta) : \mathcal{A}_\Lambda \rightarrow \mathcal{A}_\Lambda$ is given on generators as follows. Assemble the degree 1 generators c_{ij} into an $N \times N$ matrix: then*

$$\mathcal{A}(\vartheta)(C) = \tilde{P}_\gamma C \tilde{P}_\gamma^{-1}.$$

For degree 0 generators, $\mathcal{A}(\vartheta)$ fixes all non-satellite crossings, while its action on degree 0 satellite crossings is as follows:

$$\begin{aligned} \mathcal{A}(\vartheta) \begin{pmatrix} h_1 \\ \vdots \\ h_n \end{pmatrix} &= P_\gamma \begin{pmatrix} h_1 \\ \vdots \\ h_n \end{pmatrix} \\ \mathcal{A}(\vartheta) \begin{pmatrix} h'_1 \\ \vdots \\ h'_n \end{pmatrix} &= (P_\gamma^T)^{-1} \begin{pmatrix} h'_1 \\ \vdots \\ h'_n \end{pmatrix}. \end{aligned}$$

Here h_1, \dots, h_n is any group of satellite crossings coming from a crossing of β where the i -th strand is the overcrossing, while h'_1, \dots, h'_n is any group of satellite crossings coming from a crossing of β where the i -th strand is the undercrossing. See Figure 24.




FIGURE 24. A group of satellite crossings coming from an overcrossing (left) and an undercrossing (right).

Proof. The ϑ -loop consists of a sequence of Reidemeister III moves that push the purple box around, and consequently the map $\mathcal{A}(\vartheta)$ is the composition of a sequence of algebra isomorphisms corresponding to these Reidemeister III moves, as given concretely in Section 4.1.2. In particular, any non-satellite crossing does not participate in any of the Reidemeister III moves and so it is fixed by $\mathcal{A}(\vartheta)$.

Next consider a group of degree 0 satellite crossings h_1, \dots, h_n as in the statement of the proposition (the argument for h'_1, \dots, h'_n is similar and will be omitted). The ϑ -loop pushes the purple box containing γ through h_1, \dots, h_n from right to left. Since the path matrix P_γ is a product of path matrices for individual crossings and columns of base points, and we can factor the action of $\mathcal{A}(\vartheta)$ on h_1, \dots, h_n by pushing each individual crossing and base point column across h_1, \dots, h_n from right to left and composing the results, the key is to observe what happens when we push a single crossing or base point column of γ across h_1, \dots, h_n .




FIGURE 25. Pushing a crossing (left) or a column of base points (right) across a group of satellite crossings.

If we push a crossing a from γ across h_1, \dots, h_n by a single Reidemeister III move as shown in Figure 25 (left), then from Section 4.1.2, the associated isomorphism sends $a \mapsto a$, $h_i \mapsto h_{i+1}$, and $h_{i+1} \mapsto h_i + ah_{i+1}$. (Note that compared to Figure 19, the crossings h_i and h_{i+1} have switched places after the Reidemeister III move.) This is precisely the matrix map

$$\begin{pmatrix} h_i \\ h_{i+1} \end{pmatrix} \mapsto \begin{pmatrix} 0 & 1 \\ 1 & a \end{pmatrix} \begin{pmatrix} h_i \\ h_{i+1} \end{pmatrix}.$$

Thus pushing the crossing a across h_1, \dots, h_n acts on $\begin{pmatrix} h_1 \\ \vdots \\ h_n \end{pmatrix}$ by left multiplication by the path matrix for a . If instead we push a column of base points across h_1, \dots, h_n as shown in Figure 25 (right), then from Section 4.1.1, the associated isomorphism sends h_i to $t_i h_i$ for $i \in [1, n]$, which corresponds to left multiplication by the diagonal matrix with diagonal entries t_1, \dots, t_n . Composing the individual isomorphisms, we conclude that the purple-box monodromy indeed acts on h_1, \dots, h_n by left multiplication by the path matrix P_γ , as desired.




FIGURE 26. Pushing the purple box through the pigtail.

Finally, we consider the degree 1 crossings. As we perform the ϑ -loop, the purple box passes through the “pigtail” region with the degree 1 crossings twice; see Figure 26. On the first pass, when the purple box pushes through the strand containing crossings c_{1j}, \dots, c_{nj} , this yields the map that sends $\begin{pmatrix} c_{1j} \\ \vdots \\ c_{nj} \end{pmatrix} \mapsto P_\gamma \begin{pmatrix} c_{1j} \\ \vdots \\ c_{nj} \end{pmatrix}$, by the same argument as for the degree 0 crossings h_1, \dots, h_n above. Thus the first pass of the purple box through the degree-1 region cumulatively has the effect of sending C to $\tilde{P}_\gamma C$. Similarly the second pass (from upper right to lower left) sends C to $C \tilde{P}_\gamma^{-1}$. Together, the two passes send C to $\tilde{P}_\gamma C \tilde{P}_\gamma^{-1}$. \square

5.3. The Kálmán loop. The techniques of this section can be applied to compute the monodromy of other loops of Legendrian links besides ϑ -loops. One case where it is particularly simple to calculate the monodromy in our setting is the loop of Legendrian $T(p, q)$ -torus links originally studied by Kálmán in [Kál05]. For concreteness we focus here on the most basic example of the Kálmán loop, which involves the max-tb Legendrian right handed trefoil ($p = 2, q = 3$). Kálmán constructs a loop in the space of these Legendrian trefoils and proves that the induced action on the degree-0 Legendrian contact homology has order 5. Here we reinterpret this result in our setting.

In our language, the Legendrian trefoil is the (-1) -closure of the admissible 2-stranded braid σ_1^5 ; in other words, it is $\Lambda(\beta, 1; \gamma) \subset (\mathbb{R}^3, \xi_{\text{st}})$ where $\beta \in \text{Br}_1^+$ is the 1-stranded braid and $\gamma = \sigma_1^5 \in \text{Br}_2^+$. We label the crossings of γ a_1, \dots, a_5 and place base points t_1, t_2 to the right of γ , as shown in Figure 27 (left). The ϑ -loop moves the entire braid σ_1^5 around the standard unknot $\Lambda(\beta) = U$ until it returns to its starting point. We can factor this loop as the fifth power of another loop δ , which moves the single leftmost crossing of σ_1^5 around the unknot until it returns to γ as the rightmost crossing. Note that this move shifts the position of the base points t_1, t_2 ; we then slide t_1, t_2 along the knot until they return to their original positions. See Figure 27. The combination of the crossing move and the base point move forms a loop beginning and ending at $\Lambda(\beta, 1; \gamma)$, which is the Kálmán loop and which we denote by δ .




FIGURE 27. The Legendrian trefoil is the (-1) -closure of the depicted braid σ_1^5 . Left, the braiding region with base points; middle, the result of moving the leftmost crossing to the right; right, the result of sliding the base points back to their original position (with the slides shown in the middle diagram).

The action $\mathcal{A}(\delta)$ of the loop δ on the Legendrian contact DGA $\mathcal{A}(\Lambda(\beta, 1; \gamma))$ is easy to describe. Moving a_1 to the right simply permutes the a_i : $a_1 \mapsto a_2, \dots, a_4 \mapsto a_5, a_5 \mapsto a_1$. From Section 4.1.1, sliding the base points as indicated in Figure 27 fixes a_1 and sends a_i to $t_1^{-1}a_it_2$ for $i = 3, 5$ and $t_2^{-1}a_it_1$ for $i = 2, 4$. Thus $\mathcal{A}(\delta)$ acts on the DGA as follows:

$$a_1 \mapsto t_2^{-1}a_2t_1 \quad a_2 \mapsto t_1^{-1}a_3t_2 \quad a_3 \mapsto t_2^{-1}a_4t_1 \quad a_4 \mapsto t_1^{-1}a_5t_2 \quad a_5 \mapsto a_1.$$

(The degree 1 generators are fixed by $\mathcal{A}(\delta)$.) By inspection we see that $\mathcal{A}(\delta)$ has order 5, in agreement with Kálmán's result: $\mathcal{A}(\delta)^5 = \mathcal{A}(\vartheta)$ is the identity map.

This argument readily generalizes to (p, q) -torus links for arbitrary positive p, q . In the general case, the Legendrian link is the (-1) -closure of the admissible braid $(\sigma_1 \dots \sigma_{p-1})^{p+q} \in \text{Br}_p^+$, and the Kálmán loop δ moves the leftmost $p-1$ crossings around the unknot. As in the case of the trefoil, we immediately see that δ^{p+q} acts as the identity on the Legendrian contact DGA.

Remark 5.4. The original proof in [Kál05] that $\mathcal{A}(\delta)$ has order $p+q$ uses the Legendrian link given by the resolution of the rainbow closure of the braid $(\sigma_1 \dots \sigma_{p-1})^q$. The DGA for this link has $(p-1)q$ generators in degree 0, and Kálmán's computation of the monodromy of δ on this DGA is rather nontrivial, both because Reidemeister II moves are involved and because the DGA differential itself is quite complicated due to the presence of non-embedded disks. Kálmán then performs an intricate computation to show that this monodromy has order $p+q$. The mere fact that $p+q$ appears here, e.g. instead of p or q , is rather mysterious

from the geometric viewpoint. By contrast, in our setup with (-1) -closures, the DGA has $(p-1)(p+q)$ generators in degree 0, the monodromy of δ simply cyclically permutes these generators, and it is evident without computation that this action has order $p+q$. \square

5.4. The saddle cobordism map for (-1) -closures. So far we have discussed the ϑ -monodromy. We now turn to the other principal computational ingredient in our calculations for the upcoming Section 6, namely the calculation of saddle cobordism maps: we will consider augmentations corresponding to specific decomposable fillings, and these augmentations are the composition of a number of saddle maps.

In Section 4.2, we defined a combinatorial cobordism map $\Phi^{\text{comb}} : \mathcal{A}_{\Lambda_+} \rightarrow \mathcal{A}_{\Lambda_-}$ associated to a saddle cobordism at any proper contractible Reeb chord. This combinatorial formula allows us in Section 6 to calculate the augmentations corresponding to particular fillings of (-1) -closures, and the reader may skip ahead to that section at this point. In the present subsection, we take a slight detour and discuss what the formula for Φ^{comb} looks like for saddle cobordisms of (-1) -closures, in terms of the matrix formula for the DGA of a (-1) -closure from Section 5.1. In particular, this will allow us to see combinatorially that Φ^{comb} is indeed a chain map in this case, without going through the general theory. The interested reader may want to compare our discussion here with [GSW20a, section 3.3], which presents an independent but rather similar matrix treatment of saddle cobordism maps.

Consider a saddle cobordism whose top end is a Legendrian (-1) -closure $\Lambda_+ = \Lambda(\sigma_{k_1} \cdots \sigma_{k_r})$, and whose bottom end is the Legendrian link Λ_- obtained by resolving a contractible proper crossing of Λ_+ . For ease of notation, we will assume that the crossing is a_1 , corresponding to the braid generator σ_{k_1} , and so $\Lambda_- = \Lambda(\sigma_{k_2} \cdots \sigma_{k_r})$. (The case of a saddle resolving an arbitrary crossing a_ℓ is easy to deduce from this; just perform the cyclic-permutation isotopy sending $\Lambda_+ = \Lambda(\sigma_{k_1} \cdots \sigma_{k_r})$ to $\Lambda(\sigma_{k_\ell} \cdots \sigma_{k_r} \sigma_{k_1} \cdots \sigma_{k_{\ell-1}})$ and similarly for Λ_- .)

From Section 5.1 above, we can write down the differentials ∂_\pm on Λ_\pm in matrix form. Specifically, as in Section 5.1, we place base points $t_{\ell,i}$, $1 \leq i \leq n$, $1 \leq \ell \leq r$, next to the crossings of Λ_+ . Then Λ_- inherits this same array of base points, along with two new base points in place of the crossing a_1 , one on strand $k_1 + 1$ labeled by s_1 and one on strand k_1 labeled by $-s_1^{-1}$. By Proposition 5.2, in the notation from Section 5.1, the differentials ∂_+ and ∂_- for the DGAs of Λ_+ and Λ_- are given by the matrix formulas:

$$\begin{aligned} \partial_+(C) &= \mathbf{1} + P_{k_1}(a_1)D(\mathbf{t}_1)P_{k_2}(a_2)D(\mathbf{t}_2) \cdots P_{k_r}(a_r)D(\mathbf{t}_r) \\ \partial_-(C) &= \mathbf{1} + D(\mathbf{t}_0)D(\mathbf{t}_1)P_{k_2}(a_2)D(\mathbf{t}_2) \cdots P_{k_r}(a_r)D(\mathbf{t}_r) \end{aligned}$$

where $\mathbf{t}_0 = (1, \dots, 1, -s_1^{-1}, s_1, 1, \dots, 1)$ (with $-s_1^{-1}$ and s_1 in the k_1 and $k_1 + 1$ components respectively).

Let $\Phi^{\text{comb}} = \Phi^{\leftarrow} \circ \Phi^{\rightarrow} \circ \Phi^0 : \mathcal{A}_{\Lambda_+} \rightarrow \mathcal{A}_{\Lambda_-}$ be the cobordism map from Proposition 4.8. We first note that the action of Φ^{comb} on degree-1 Reeb chords c_{ij} is easy to write down. Indeed, write $T_{k_1}^{\leftarrow}(s_1)$ for the $n \times n$ matrix equal to the identity matrix except with $(k_1, k_1 + 1)$ entry given by s_1^{-1} . Then we have

$$(5.1) \quad \Phi^{\text{comb}}(C) = T_{k_1}^{\leftarrow}(s_1)C(T_{k_1}^{\leftarrow}(s_1))^{-1}.$$

This can be seen directly from an inspection of Figure 23, using the fact that $\Delta_{a_1}^{\rightarrow}(c_{ij}) = \emptyset$, while the only possible disks in $\Delta_{a_1}^{\leftarrow}(c_{ij})$ are thin disks heading left from their $+$ corner at a_1 , following the figure eight, and ending in the region containing the c_{ij} 's. We omit the details here.

The explicit nature of this algebraic model allows us to sketch a direct argument for why Φ^{comb} is a chain map. Note that this argument is mainly provided for context and is not

needed in the rest of the paper,¹⁶ and so we do not provide full details; see also [GSW20a, section 3.3] for a related discussion with more details.

Proposition 5.5. $\Phi^{\text{comb}} \circ \partial_+ = \partial_- \circ \Phi^{\text{comb}}$.

Proof. In order to show that Φ^{comb} is a chain map, it suffices to show that $\Phi^{\text{comb}}(\partial_+(C)) = \partial_-(\Phi^{\text{comb}}(C))$. Note that

$$\Phi^{\text{comb}}(P_{k_1}(a_1)) = P_{k_1}(s_1) = T_{k_1}^{\leftarrow}(s_1)D(\mathbf{t}_0)T_{k_1}^{\rightarrow}(s_1)$$

where $T_{k_1}^{\rightarrow}(s_1)$ is the identity matrix except with $(k_1 + 1, k_1)$ entry given by s_1^{-1} . Since Φ acts on C by conjugation by $T_{k_1}^{\leftarrow}(s_1)$, showing that Φ is a chain map reduces to verifying the following:

$$(5.2) \quad \Phi^{\text{comb}}(T_{k_1}^{\rightarrow}(s_1)D(\mathbf{t}_1)P_{k_2}(a_2)D(\mathbf{t}_2) \cdots P_{k_r}(a_r)D(\mathbf{t}_r)T_{k_1}^{\leftarrow}(s_1)) = D(\mathbf{t}_1)P_{k_2}(a_2)D(\mathbf{t}_2) \cdots P_{k_r}(a_r)D(\mathbf{t}_r).$$

Call a matrix *lower-unipotent* if it is of the form $\mathbf{1} + N$ where N is a strictly lower triangular matrix; that is, a lower-unipotent matrix is a lower triangular matrix with 1's along the diagonal. Note in particular that $T_{k_1}^{\rightarrow}(s_1)$ is lower-unipotent. Next we observe that if T is lower-unipotent then

$$T' = (P_{k_\ell}(a_\ell + T_{k_\ell+1, k_\ell}))^{-1} T P_{k_\ell}(a_\ell)$$

is again lower-unipotent: this follows from the identity of 2×2 matrices

$$\begin{pmatrix} 0 & 1 \\ 1 & a_\ell + T_{k_\ell+1, k_\ell} \end{pmatrix}^{-1} \begin{pmatrix} 1 & 0 \\ T_{k_\ell+1, k_\ell} & 1 \end{pmatrix} \begin{pmatrix} 0 & 1 \\ 1 & a_\ell \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

We can thus inductively define a sequence of lower-unipotent matrices $T'_1, T_1, T'_2, T_2, \dots, T'_r, T_r$ as follows:

$$\begin{aligned} T'_1 &= T_{k_1}^{\rightarrow}(s_1), \\ T_\ell &= D(\mathbf{t}_\ell)^{-1} T'_\ell D(\mathbf{t}_\ell), \\ T'_\ell &= (P_{k_\ell}(a_\ell + (T_{\ell-1})_{k_\ell+1, k_\ell}))^{-1} T P_{k_\ell}(a_\ell). \end{aligned}$$

Then we have

$$T_{\ell-1} P_{k_\ell}(a_\ell) D(\mathbf{t}_\ell) = P_{k_\ell}(a_\ell + (T_{\ell-1})_{k_\ell+1, k_\ell}) D(\mathbf{t}_\ell) T_\ell.$$

Write $x_\ell := (T_{\ell-1})_{k_\ell+1, k_\ell}$ for short; we now have

$$\begin{aligned} T_{k_1}^{\rightarrow}(s_1) D(\mathbf{t}_1) P_{k_2}(a_2) D(\mathbf{t}_2) \cdots P_{k_r}(a_r) D(\mathbf{t}_r) \\ = D(\mathbf{t}_1) P_{k_2}(a_2 + x_2) D(\mathbf{t}_2) P_{k_3}(a_3 + x_3) \cdots P_{k_r}(a_r + x_r) D(\mathbf{t}_r) T_r. \end{aligned}$$

The key fact now, whose proof (and precise statement) we omit here, is that the matrices T_ℓ have geometric meaning: for $i > j$, the (i, j) entry in T_ℓ counts embedded disks whose leftmost end is a positive corner at a_1 and whose rightmost end is a vertical line segment connecting strands i and j just to the right of crossing a_ℓ . (In particular, $T_r = \mathbf{1}$.) Furthermore, the map Φ^{\rightarrow} from Section 4.2 is constructed exactly to satisfy

$$\Phi^{\rightarrow}(a_\ell + x_\ell) = a_\ell$$

for all $\ell = 2, \dots, r$. As a consequence, we have

$$\Phi^{\rightarrow}(T_{k_1}^{\rightarrow}(s_1) D(\mathbf{t}_1) P_{k_2}(a_2) D(\mathbf{t}_2) \cdots P_{k_r}(a_r) D(\mathbf{t}_r)) = D(\mathbf{t}_1) P_{k_2}(a_2) D(\mathbf{t}_2) P_{k_3}(a_3) \cdots P_{k_r}(a_r) D(\mathbf{t}_r).$$

Similarly, Φ^{\leftarrow} satisfies

$$\Phi^{\leftarrow}(D(\mathbf{t}_1) P_{k_2}(a_2) D(\mathbf{t}_2) \cdots P_{k_r}(a_r) D(\mathbf{t}_r) T_{k_1}^{\leftarrow}(s_1)) = D(\mathbf{t}_1) P_{k_2}(a_2) D(\mathbf{t}_2) P_{k_3}(a_3) \cdots P_{k_r}(a_r) D(\mathbf{t}_r).$$

Combining this equation and the previous equation now yields (5.2), whence Φ^{comb} is a chain map. \square

¹⁶The computation in the proof of Proposition 5.5 does contribute to the implementation of the program [Ng], in the code calculating the augmentation associated to a filling of a (-1) -closure.

6. PROOF OF INFINITELY MANY FILLINGS

In this section we prove Theorem 1.1. First, we describe the scheme of proof that we will use for all the Legendrian links $\Lambda \in \mathcal{H}$. The cases $\Lambda(\tilde{D}_4), \Lambda_1, \Lambda_2$ and $\Lambda(\beta_{11}), \Lambda(\beta_{12}), \Lambda(\beta_{21})$ are then proven directly using this strategy. The general cases $\Lambda(\tilde{D}_n), \Lambda_n$ are concluded from Proposition 6.5 and the proofs we give for the two cases $\Lambda(\tilde{D}_4), \Lambda_2$.

6.1. The argument. Let $\Lambda \subset (\mathbb{R}^3, \xi_{\text{st}})$ be a Legendrian link $\Lambda = \Lambda(\beta, i; \gamma)$, $\beta \in \text{Br}_N^+, \gamma \in \text{Br}_M^+$, and consider its ϑ -loop, as introduced in Section 2.4. The general structure of our proofs can be described in three steps, as follows:

- (i) First, choose an ordered sequence of crossings for β and γ such that resolving these crossings yields an orientable exact Lagrangian filling $L \subset (\mathbb{R}^4, \lambda_{\text{st}})$ of the Legendrian link Λ .
- (ii) Second, compute the augmentation $\varepsilon_L : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}[H_1(L) \oplus \mathbb{Z}^{m-1}]$ associated to the exact Lagrangian filling L (where m is the number of components of Λ) and the induced maps $\vartheta^k : \mathcal{A}_\Lambda \rightarrow \mathcal{A}_\Lambda$, $k \in \mathbb{N}$. We note that all crossings chosen in (i) will have the property that their complement contains a half-twist, and consequently they are contractible and proper by Proposition 4.5. Thus we may apply the combinatorial formulas from Section 4.2 in this step.
- (iii) Third, fix a crossing a for the braid word β associated to the Legendrian link Λ , which we consider as one of the generators $a \in \mathcal{A}_\Lambda$ of the Legendrian contact DGA. Consider the invariant

$$E(k, a) := \max_{\eta: R \rightarrow \mathbb{Z}} |(\eta \circ \varepsilon_L \circ \vartheta^k)(a)|, \quad k \in \mathbb{N},$$

where $R = \mathbb{Z}[H_1(L) \oplus \mathbb{Z}^{m-1}]$ and $\eta : R \rightarrow \mathbb{Z}$ runs over all possible unital ring morphisms. Note that the set of such morphisms is finite, as the first Betti number $b_1(L)$ is finite, and thus $E(k, a)$ is a well-defined maximum over a finite set of integers. Finally, show that $E(k, a)$ is a strictly increasing function of $k \in \mathbb{N}$.

The different choices for the Lagrangian filling (and thus the augmentation ε_L) and crossing $a \in \mathcal{A}_\Lambda$ influence the computation of the invariant $E(k, a)$. Finding the maximum over a set whose cardinality grows exponentially in $l(\beta) + l(\gamma)$ makes brute force computation a difficult (though not unfeasible) route. Thus, particular care must be devoted in choosing the augmentation ε_L and the crossing $a \in \mathcal{A}_\Lambda$: we will find crossings $a \in \mathcal{A}_\Lambda$ and Lagrangian fillings whose augmentations satisfy that $(\varepsilon_L \circ \vartheta^k)(a)$ is a *positive* Laurent polynomial in $\mathbb{Z}[H_1(L)]$, for all $k \in \mathbb{N}$, making the invariant $E(k, a)$ readily computable.

Remark 6.1. Executing the argument laid out here for specific Legendrian links, including all of the ones that we consider in this section, is readily amenable to calculation by computer. A *Mathematica* notebook that performs the calculations contained in the remainder of this section, and is suitable for calculations for general (-1) -closures, is available at the second author's web page [Ng]. We will work out the argument for \tilde{D}_4 in detail, without recourse to the computer program, in Section 6.2 below; we provide fewer details for subsequent computations and refer the reader to the program. \square

6.2. Augmentations for $\Lambda(\tilde{D}_4)$. We now turn to proving Theorem 1.1 for the Legendrian link $\Lambda(\tilde{D}_n)$, $n \geq 4$. In this subsection we present the argument for $n = 4$; the general $n \geq 4$ case is deduced from this in Section 6.5.

As stated in the introduction, the Legendrian link $\Lambda(\tilde{D}_4) \subset (\mathbb{R}^3, \xi_{\text{st}}) \subset (\mathbb{S}^3, \xi_{\text{st}})$ is defined to be the rainbow closure of the positive braid $(\sigma_2 \sigma_1 \sigma_3 \sigma_2)^2$, which is also the (-1) -closure of

the braid $(\sigma_2\sigma_1\sigma_3\sigma_2)^4\sigma_3^2\sigma_1^2 = (\sigma_2\sigma_1\sigma_3\sigma_2)^2\Delta_4^2$. A Lagrangian projection for $\Lambda(\tilde{D}_4)$ is depicted in Figure 28. Let us prove the following result:

Theorem 6.2 (The \tilde{D}_4 -Legendrian). *Let $\vartheta : \mathbb{S}^1 \rightarrow \mathcal{L}(\Lambda(\tilde{D}_4))$ be the purple-box Legendrian loop. Then there exists a Lagrangian filling $L \subset \mathbb{D}^4$ of $\Lambda(\tilde{D}_4)$ such that the ϑ -orbit of the system of augmentations ε_L is entire.*

In order to prove Theorem 6.2, we set some notation and lay out the pieces that go into the proof. Let us label the crossings of the positive braid $(\sigma_2\sigma_1\sigma_3\sigma_2)^4\sigma_3^2\sigma_1^2$ from left to right as

$$a_{16}, \dots, a_1, a_{17}, a_{18}, a_{19}, a_{20}.$$

Figure 28 shows the Lagrangian projection of $\Lambda(\tilde{D}_4)$ that we use for the proof, where the labeled crossings are also depicted. These crossings are the degree-0 Reeb chords of a Legendrian front for $\Lambda(\tilde{D}_4)$. The ϑ -monodromy is obtained by carrying around the purple box containing the two crossings a_{19}, a_{20} and the two base points t_1, t_2 , as shown in Figure 28, cf. Figure 2.

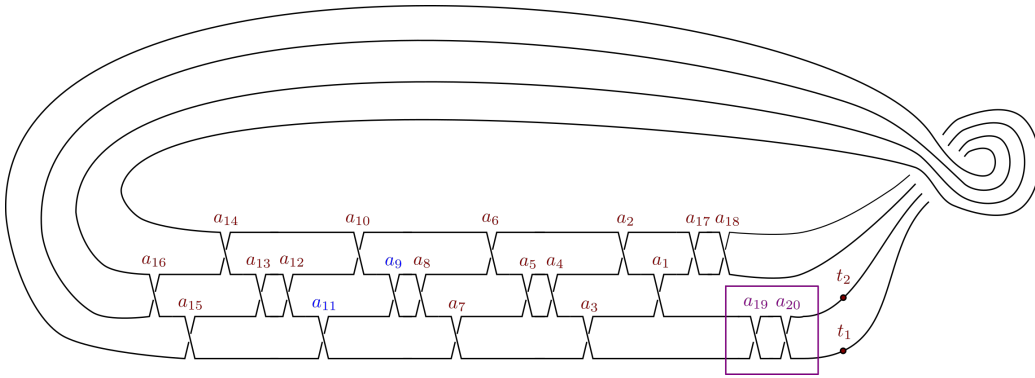


FIGURE 28. Lagrangian projection for the Legendrian link $\Lambda(\tilde{D}_4)$, as used in the proof of Theorem 6.2. The crossings a_{11}, a_9 , in blue, are used to detect the infinite order of the ϑ -monodromy. In this case, the ϑ -monodromy is obtained by moving the crossings a_{19}, a_{20} around this projection.

The filling L of $\Lambda(\tilde{D}_4)$ that we will consider is the decomposable filling constructed as follows. Resolve the following crossings of $\Lambda(\tilde{D}_4)$ in order:

$$a_9, a_{10}, a_{11}, a_{12}, a_{13}, a_{14}, a_{15}, a_{16}.$$

Note that at each step the remaining braid is admissible in the sense of Definition 2.5: this follows from Proposition 2.7 and the fact that the crossings $a_4, a_3, a_2, a_1, a_{17}, a_{19}$ comprise a half-twist. Thus each step produces a legal Lagrangian projection of a Legendrian link, and each resolved crossing is contractible. The result of resolving these 8 crossings is the (-1) -closure of a full positive twist, which we write as Λ_0 and is precisely the standard 4-component Legendrian unlink. We then fill in each of the 4 component unknots. This gives the desired filling L of $\Lambda(\tilde{D}_4)$, expressed as 8 saddle cobordisms and 4 minimum cobordisms.

Following the discussion in Section 3.5, we use $\mathbb{Z}[t_1^{\pm 1}, \dots, t_4^{\pm 1}, s_9^{\pm 1}, \dots, s_{16}^{\pm 1}]$ as the coefficient ring for the DGA $\mathcal{A}(\Lambda(\tilde{D}_4))$. We will need two maps on $\mathcal{A}(\Lambda(\tilde{D}_4))$, induced by the ϑ -monodromy and the filling L . The former map is an automorphism $\vartheta : \mathcal{A}(\Lambda(\tilde{D}_4)) \rightarrow \mathcal{A}(\Lambda(\tilde{D}_4))$. For the latter, as in Section 3.5, L induces an augmentation $\varepsilon_L : \mathcal{A}(\Lambda(\tilde{D}_4)) \rightarrow R$. Here $R = \mathbb{Z}[t_1^{\pm 1}, \dots, t_4^{\pm 1}, s_9^{\pm 1}, \dots, s_{16}^{\pm 1}]/(w_1 = w_2 = w_3 = w_4 = -1)$, where w_1, w_2, w_3, w_4 are the product of the labels of the base points on each unknot in Λ_0 . Since by inspection

t_1, t_2, t_3, t_4 appear on all distinct components of Λ_0 , the quotient allows us to solve for the t_i 's, and we conclude that $R \cong \mathbb{Z}[s_9^{\pm 1}, \dots, s_{16}^{\pm 1}]$.

Our aim is to pairwise distinguish the iterates $\varepsilon_L \circ \vartheta^k : \mathcal{A}(\Lambda(\tilde{D}_4)) \rightarrow R$, $k \in \mathbb{N}$, even up to automorphisms of R . We will do this by computing the image of a_9 under each of these maps. In order to perform this computation, we need to partially compute the maps ϑ and ε_L .

We first consider the monodromy automorphism ϑ , which we compute using Proposition 5.3. First note that ϑ fixes the variables a_{19}, a_{20}, t_1, t_2 that appear inside the purple box. We will be interested in what ϑ does to the two Reeb chords a_9, a_{11} , which are depicted in blue in Figure 28. The path matrix associated to the purple box is given by

$$M = \begin{pmatrix} 0 & 1 \\ 1 & a_{19} \end{pmatrix} \cdot \begin{pmatrix} 0 & 1 \\ 1 & a_{20} \end{pmatrix} \cdot \begin{pmatrix} t_1 & 0 \\ 0 & t_2 \end{pmatrix} = \begin{pmatrix} t_1 & t_2 a_{20} \\ t_1 a_{19} & t_2(1 + a_{19} a_{20}) \end{pmatrix}.$$

By Proposition 5.3, the effect of the DGA automorphism $\vartheta \in \text{Aut}(\mathcal{A}(\Lambda(\tilde{D}_4)))$ on the two crossings a_{11}, a_9 , which are depicted in blue in Figure 28, is

$$\begin{pmatrix} a_{11} \\ a_9 \end{pmatrix} \mapsto \vartheta \begin{pmatrix} a_{11} \\ a_9 \end{pmatrix} = M \begin{pmatrix} a_{11} \\ a_9 \end{pmatrix}.$$

Next consider the augmentation ε_L , which we can explicitly compute using the formulas from Section 4. We will only need the following partial computation:

Lemma 6.3. *We have $\varepsilon_L(a_9) = s_9$, $\varepsilon_L(a_{11}) = s_{11}$, and*

$$\begin{aligned} \varepsilon_L(t_1) &= -s_{11}s_{15}, \\ \varepsilon_L(t_2) &= -\frac{s_9 s_{12} s_{13} s_{16}}{s_{11} s_{15}}, \\ \varepsilon_L(a_{19}) &= \frac{s_9}{s_{11}} - \frac{s_9 s_{12} s_{13}}{s_{11}^2 s_{15}}, \\ \varepsilon_L(a_{20}) &= -\frac{s_{11}^3 s_{15}^2}{s_9^2 s_{12} s_{13} s_{16}} + \frac{s_{11}^2 s_{15}}{s_9 s_{12} s_{13}} - \frac{s_{11}^2 s_{15}^2}{s_9 s_{10} s_{13} s_{16}} + \frac{s_{11}^2 s_{14} s_{15}^2}{s_9 s_{12} s_{13}^2 s_{16}}. \end{aligned}$$

Proof. For $i = 9, \dots, 16$, let $\Phi_i = \Phi_{L_{a_i}}^{\text{comb}}$ denote the combinatorial cobordism map associated to the saddle cobordism at a_i , as described in Section 4.2; also let $\varepsilon_0 : \mathcal{A}_{\Lambda_0} \rightarrow R \cong \mathbb{Z}[s_9^{\pm 1}, \dots, s_{16}^{\pm 1}]$ denote the augmentation associated to the disk filling of Λ_0 . We have

$$\varepsilon_L = \varepsilon_0 \circ \Phi_{16} \circ \dots \circ \Phi_9.$$

We begin by computing ε_0 . Note that all Reeb chords of Λ_0 either have degree 1 (for the 4^2 crossings on the right) or connect different components of Λ_0 (for the crossings a_i for $1 \leq i \leq 8$ and $17 \leq i \leq 20$). Since the filling of Λ_0 consists of four disjoint disks, it follows that ε_0 sends all Reeb chords to 0. As for the t_i parameters, an inspection of Figure 28 yields that the unknot components of Λ_0 containing t_1 and t_2 contain the following base points in order: $-s_{15}^{-1}, -s_{11}^{-1}, t_1$ and $-s_{16}^{-1}, s_{15}, -s_{13}^{-1}, -s_{12}^{-1}, s_{11}, -s_9^{-1}, t_2$, respectively. Setting each of the products of these base points equal to -1 gives $t_1 = -s_{11}s_{15}$ and $t_2 = -\frac{s_9 s_{12} s_{13} s_{16}}{s_{11} s_{15}}$, and these are the respective images of t_1 and t_2 under ε_0 (and thus under ε_L as well).

We now proceed to compute ε_L for $a_9, a_{11}, a_{19}, a_{20}$. The sequence of saddle moves has been chosen to simplify the computation of $\varepsilon_L(a_9)$ and $\varepsilon_L(a_{11})$: indeed, $\varepsilon_L(a_9) = \Phi_9(a_9) = s_9$, while Φ_9 and Φ_{10} fix a_{11} and so $\varepsilon_L(a_{11}) = \Phi_{11}(a_{11}) = s_{11}$.

For a_{19} , we keep track of disks with two positive punctures, one at a_{19} and one at the crossing being resolved. There are no such disks when we resolve a_9 and a_{10} . When we resolve a_{11} ,




FIGURE 29. Disks with positive corners at a_{11}, a_{19} (left) and a_{15}, a_{19} (right), contributing to $\Phi_{11}(a_{19})$ and $\Phi_{15}(a_{19})$ respectively. For these disks, the positive corner at a_{19} has positive orientation sign, while the positive corner at a_{11} and a_{15} has negative orientation sign.

there is one disk $\Delta \in \Delta_{a_{11}}^{\rightarrow}(a_{19})$ passing through $-s_9^{-1}$ (against the orientation) with no negative corners; see Figure 29. From Definition 4.6, we read off $\Phi^{\rightarrow}(a_{19}) = a_{19} + s_9 s_{11}^{-1}$ and $\Phi^{\leftarrow}(a_{19}) = a_{19}$, and so $\Phi_{11}(a_{19}) = a_{19} + s_9 s_{11}^{-1}$. As we successively resolve a_{12}, \dots, a_{16} , the only additional relevant disk with two positive punctures comes when we resolve a_{15} and is shown in Figure 29; this gives $\Phi_{15}(a_{19}) = a_{19} - s_9 s_{11}^{-1} s_{12} s_{13} s_{15}^{-1} s_{11}^{-1}$. We conclude that

$$\varepsilon_L(a_{19}) = \varepsilon_0(\Phi_{15}(\Phi_{11}(a_{19}))) = \varepsilon_0\left(a_{19} + \frac{s_9}{s_{11}} - \frac{s_9 s_{12} s_{13}}{s_{11}^2 s_{15}}\right) = \frac{s_9}{s_{11}} - \frac{s_9 s_{12} s_{13}}{s_{11}^2 s_{15}}.$$

The computation of $\varepsilon_L(a_{20})$ is similar but slightly more involved. We compute that

$$\begin{aligned}\Phi_9(a_{20}) &= a_{20} - t_1 a_{11} s_9^{-1} t_2^{-1} \\ \Phi_{10}(a_{20}) &= a_{20} - t_1 a_{12} s_{10}^{-1} t_2^{-1} \\ \Phi_{13}(a_{20}) &= a_{20} + t_1 s_{13}^{-1} a_{14} t_2^{-1} \\ \Phi_{15}(a_{20}) &= a_{20} + t_1 s_{15}^{-1} a_{16} t_2^{-1};\end{aligned}$$

piecing these together, along with $\Phi_i(a_i) = s_i$ and the values computed above for $\varepsilon_0(t_1)$ and $\varepsilon_0(t_2)$, gives the desired expression for $\varepsilon_L(a_{20})$. \square

We are now in position to prove Theorem 6.2.

Proof of Theorem 6.2. Consider the following matrices with entries in $\mathbb{Z}[s_9^{\pm 1}, \dots, s_{16}^{\pm 1}]$:

$$M_0 := \varepsilon_L(M), \quad v_0 := \varepsilon_L\begin{pmatrix} a_{11} \\ a_9 \end{pmatrix} = \begin{pmatrix} s_{11} \\ s_9 \end{pmatrix}, \quad N := \begin{pmatrix} s_{11} & 1 \\ s_9 & 0 \end{pmatrix}, \quad M_1 := N^{-1} M_0 N.$$

For $k \in \mathbb{N}$, the augmentation $\varepsilon_L \circ \vartheta^k$ sends the column vector $\begin{pmatrix} a_{11} \\ a_9 \end{pmatrix}$ to

$$\varepsilon_L(M)^k \cdot \varepsilon_L\begin{pmatrix} a_{11} \\ a_9 \end{pmatrix} = M_0^k v_0 = N(N^{-1} M_0 N)^k \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix} = N M_1^k \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

We can explicitly write down M_1 using Lemma 6.3. This leads to the following observation: if we replace $s_{11}, s_{12}, s_{15}, s_{16}$ by their negatives $-s_{11}, -s_{12}, -s_{15}, -s_{16}$, the matrix M_1 becomes

$$(M_1)|_{\{s_j \rightarrow -s_j, j=11,12,15,16\}} = \begin{pmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{pmatrix},$$

where the entries are

$$\begin{aligned}m_{11} &= \frac{s_9 s_{13} s_{12}^2}{s_{10} s_{11}} + \frac{s_9 s_{14} s_{12}}{s_{11}} + \frac{s_9 s_{15} s_{12}}{s_{10}} + \frac{s_9 s_{14} s_{15}}{s_{13}} + s_9 s_{16} \\ m_{12} &= \frac{s_{12} s_{13}}{s_{11}} + s_{15} \\ m_{21} &= \frac{s_9 s_{13} s_{12}^2}{s_{10}} + s_9 s_{14} s_{12} \\ m_{22} &= s_{12} s_{13}.\end{aligned}$$

Note that all the coefficients are *positive* Laurent polynomials in the variables s_9, \dots, s_{16} : this is the algebraic reason for our choice of augmentation ε_L , and the change of signs for the variables $s_{11}, s_{12}, s_{15}, s_{16}$.

Let us now finally conclude that the iterates $\varepsilon_L \circ \vartheta^k$ are pairwise distinct. We do this by studying the quantity

$$E(k, a_9) := \max_{\eta: R \rightarrow \mathbb{Z}} |(\eta \circ \varepsilon_L \circ \vartheta^k)(a_9)|,$$

where $\eta : R \rightarrow \mathbb{Z}$ runs over all possible 2^8 unital ring morphisms. This is an integer-valued invariant of an augmentation $\varepsilon_L : \mathcal{A}_\Lambda \rightarrow R$ even up to post-composition of an automorphism of R . That is, if $E(k, a_9) \neq E(l, a_9)$ then there exists no automorphism $\varphi \in \text{Aut}(R)$ such that $\varphi(\varepsilon_L \circ \vartheta^k) = \varepsilon_L \circ \vartheta^l$, and thus the k -th and l -th ϑ -iterates of ε_L are distinct. In order to compute $E(k, a_9)$, we note that

$$|(\varepsilon_L \circ \vartheta^k)(a_9)| = \left| (s_9 \ 0) M_1^k \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right| = \left| (1 \ 0) M_1^k \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right|$$

is the absolute value of the upper-left entry of M_1^k . A unital ring morphism $\eta : R \rightarrow \mathbb{Z}$ is uniquely determined by specifying the values $s_9, \dots, s_{16} \in \{\pm 1\}$ and since the entries $m_{11}, m_{12}, m_{21}, m_{22}$ are positive Laurent polynomials, the value $|(\varepsilon_L \circ \vartheta^k)(a_9)|$ is maximized when $s_i = -1$ for $i = 11, 12, 15, 16$ and $s_i = 1$ for $i = 9, 10, 13, 14$. It follows that $E(k, a_9)$ is equal to the upper-left entry of $(\begin{smallmatrix} 5 & 2 \\ 2 & 1 \end{smallmatrix})^k$, which is a strictly increasing function of k . This proves that $E(k, a_9) \neq E(l, a_9)$ if $k \neq l$, as required. \square

6.3. Three Variations on the Affine D_4 -braid. Let us next consider the following three Legendrian links from the Introduction:

$$\begin{aligned} \Lambda(\beta_{12}) &= \Lambda((\sigma_1 \sigma_2 \sigma_2 \sigma_1)^2 \sigma_1, 1; \sigma_1^2), & \Lambda(\beta_{21}) &= \Lambda((\sigma_1 \sigma_2 \sigma_2 \sigma_1)^2 \sigma_1^2, 1; \sigma_1), \\ \Lambda(\beta_{11}) &= \Lambda((\sigma_1 \sigma_2 \sigma_2 \sigma_1)^2 \sigma_1, 1; \sigma_1). \end{aligned}$$

These are obtained from the \tilde{D}_4 -braid by removing the crossing a_{18} , for β_{12} , the crossing a_{20} , for β_{21} or the two crossings a_{18}, a_{20} , for the braid β_{11} . See Figure 28 for the notation on the crossings, we denote the crossings of these three braids by the same labels¹⁷ as in Figure 28. In these three cases, we can use the template given by the proof of Theorem 6.2, again by studying the crossings a_9, a_{11} . We will omit the details and just give the choice of Lagrangian filling L , its corresponding augmentation ε_L as computed from the formulas in Section 4.2, and the augmented matrices M_1 . These computations are also contained in the *Mathematica* notebook [Ng].

- **The link $\Lambda(\beta_{12})$.** The Lagrangian filling L is obtained by resolving the crossings $a_9, a_{10}, a_{11}, a_{12}, a_{13}, a_{15}, a_{16}$ in order. The augmentation ε_L sends

$$\begin{aligned} t_1 &\rightarrow -s_{11}s_{15}, t_2 \rightarrow -\frac{s_9 s_{12} s_{13} s_{16}}{s_{11} s_{15}}, a_9 \rightarrow s_9, a_{11} \rightarrow s_{11}, a_{19} \rightarrow \frac{s_9}{s_{11}} + \frac{s_{12} s_{13} s_9}{s_{11}^2 s_{15}}, \\ a_{20} &\rightarrow -\frac{s_{15}^2 s_{11}^3}{s_9^2 s_{12} s_{13} s_{16}} - \frac{s_{15} s_{11}^2}{s_9 s_{12} s_{13}} - \frac{s_{15}^2 s_{11}^2}{s_9 s_{10} s_{13} s_{16}}. \end{aligned}$$

The augmented matrix $M_1 = N^{-1} M_0 N$ satisfies

$$M_1|_{\{s_j \rightarrow -s_j, j=11,12,15,16\}} = \begin{pmatrix} \frac{s_9 s_{13} s_{12}^2}{s_{10} s_{11}} + \frac{s_9 s_{15} s_{12}}{s_{10}} + s_9 s_{16} & \frac{s_{12} s_{13}}{s_{11}} + s_{15} \\ \frac{s_9 s_{12}^2 s_{13}}{s_{10}} & s_{12} s_{13} \end{pmatrix},$$

whose entries are all positive Laurent polynomials.

¹⁷That is, the crossing a_i for the \tilde{D}_4 -braid is still denoted a_i for the braids β_{ij} , $1 \leq i, j \leq 2$, where β_{22} is precisely the \tilde{D}_4 -braid.

- **The link $\Lambda(\beta_{21})$.** The Lagrangian filling L is obtained by resolving the crossings $a_9, a_{10}, a_{11}, a_{12}, a_{13}, a_{14}, a_{16}$ in order. This augmentation ε_L sends

$$t_1 \rightarrow \frac{s_9 s_{12} s_{13}}{s_{11}}, t_2 \rightarrow -s_{11} s_{16}, a_9 \rightarrow s_9, a_{11} \rightarrow s_{11}, a_{19} \rightarrow \frac{s_9 s_{13} s_{12}^2}{s_{10} s_{11}^2 s_{16}} + \frac{s_{13} s_{12}}{s_{11} s_{16}} - \frac{s_9 s_{14} s_{12}}{s_{11}^2 s_{16}} + \frac{s_9}{s_{11}}.$$

The augmented matrix $M_1 = N^{-1} M_0 N$ satisfies

$$M_1|_{\{s_j \rightarrow -s_j, j=11,12,16\}} = \begin{pmatrix} \frac{s_9 s_{13} s_{12}^2}{s_{10} s_{11}} + \frac{s_9 s_{14} s_{12}}{s_{11}} + s_9 s_{16} & \frac{s_{12} s_{13}}{s_{11}} \\ \frac{s_9 s_{13} s_{12}^2}{s_{10}} + s_9 s_{14} s_{12} & s_{12} s_{13} \end{pmatrix},$$

whose entries are all positive Laurent polynomials.

- **The link $\Lambda(\beta_{11})$.** The Lagrangian filling L is obtained by resolving the crossings $a_9, a_{10}, a_{11}, a_{12}, a_{13}, a_{16}$ in order. This augmentation ε_L sends

$$t_1 \rightarrow \frac{s_9 s_{12} s_{13}}{s_{11}}, t_2 \rightarrow -s_{11} s_{16}, a_9 \rightarrow s_9, a_{11} \rightarrow s_{11}, a_{19} \rightarrow \frac{s_9 s_{13} s_{12}^2}{s_{10} s_{11}^2 s_{16}} + \frac{s_{13} s_{12}}{s_{11} s_{16}} + \frac{s_9}{s_{11}}.$$

The augmented matrix $M_1 = N^{-1} M_0 N$ satisfies

$$M_1|_{\{s_j \rightarrow -s_j, j=11,12,16\}} = \begin{pmatrix} \frac{s_9 s_{13} s_{12}^2}{s_{10} s_{11}} + s_9 s_{16} & \frac{s_{12} s_{13}}{s_{11}} \\ \frac{s_9 s_{12}^2 s_{13}}{s_{10}} & s_{12} s_{13} \end{pmatrix},$$

whose entries are all positive Laurent polynomials.

This completes the proof of Theorem 1.1 for the Legendrian links $\Lambda(\beta_{11}), \Lambda(\beta_{12}), \Lambda(\beta_{21})$. We emphasize that these links all have a stabilized component (or two, in the case of $\Lambda(\beta_{11})$). In particular, these Legendrian links are not the rainbow closure of a positive braid, and our Floer-theoretic argument is presently the only known argument that shows the existence of infinitely many Lagrangian fillings for these Legendrian links.




FIGURE 30. Lagrangian projection for the Legendrian link Λ_1 , as used in the proof of Theorem 6.4. In this proof, the crossings a_{10}, a_9 , in blue, are used to detect the infinite order of the ϑ -monodromy. The ϑ -monodromy is obtained by moving the purple box around this projection. To construct Λ_2 instead, add an additional crossing labeled a_{14} to the purple box, between a_{13} and the base points t_1, t_2 .

6.4. Monodromy for the Braids Λ_1 and Λ_2 . We next prove Theorem 1.1 for the Legendrian links Λ_1 and Λ_2 from the Introduction. Figure 30 depicts a Lagrangian projection of Λ_1 ; Λ_2 comes from adding one additional crossing to the purple box.

Theorem 6.4 (The Λ_1 - and Λ_2 -Legendrians). *Let Λ_1 and Λ_2 be the (-1) -closures of the 3-braids $(\sigma_2\sigma_1\sigma_1\sigma_2)^3\sigma_1$ and $(\sigma_2\sigma_1\sigma_1\sigma_2)^3\sigma_1^2$, respectively. Let $\vartheta : \mathbb{S}^1 \rightarrow \mathcal{L}(\Lambda_1)$ be the purple-box Legendrian loop. Then for $n = 1, 2$, there exists a Lagrangian filling $L \subset \mathbb{D}^4$ of Λ_n such that the ϑ -orbit of the system of augmentations ε_L is entire.*

Proof. As in Section 6.3, this follows the proof of Theorem 6.2 and we will simply specify the fillings and describe the corresponding augmentations and augmented matrices. The computation of the augmentations can be found in [Ng].

We begin with Λ_1 , whose Lagrangian projection is shown in Figure 30. We choose the filling L given by resolving the crossings

$$a_{10}, a_9, a_8, a_7, a_6, a_5, a_{12}$$

in order. (As usual, this produces an unlink, and we then fill in each of the 3 unknot components to complete the construction of L .) The augmentation ε_L sends

$$\begin{aligned} t_1 &\rightarrow \frac{s_5 s_8 s_9}{s_6 s_7 s_{10}}, t_2 \rightarrow -s_6 s_7 s_{10} s_{12}, a_9 \rightarrow s_9, a_{10} \rightarrow s_{10}, \\ a_{13} &\rightarrow \frac{s_5}{s_6} - \frac{s_7}{s_8} + \frac{s_9}{s_{10}} + \frac{s_5 s_8}{s_6^2 s_7^2 s_{10} s_{12}} - \frac{s_5}{s_6^2 s_7^2 s_9 s_{10} s_{12}} + \frac{1}{s_6 s_7 s_8 s_9 s_{10} s_{12}}. \end{aligned}$$

Define $M_0 = \varepsilon_L \begin{pmatrix} 0 & t_2 \\ t_1 & a_{13} t_2 \end{pmatrix}$, $N = \varepsilon_L \begin{pmatrix} a_{10} & 1 \\ a_9 & 0 \end{pmatrix}$, and $M_1 = N^{-1} M_0 N$; then

$$M_1|_{\{s_j \rightarrow -s_j, j=5,7,8,10\}} = \begin{pmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{pmatrix},$$

and the entries

$$\begin{aligned} m_{11} &= \frac{s_6 s_{10} s_{12} s_7^2}{s_8} + s_6 s_9 s_{12} s_7 + s_5 s_{10} s_{12} s_7 + \frac{s_5}{s_6 s_7 s_9} + \frac{1}{s_8 s_9} \\ m_{12} &= \frac{s_5 s_8}{s_6 s_7 s_{10}} \\ m_{21} &= s_5 s_7 s_{12} s_{10}^2 + \frac{s_6 s_7^2 s_{12} s_{10}^2}{s_8} + \frac{s_5 s_{10}}{s_6 s_7 s_9} + \frac{s_{10}}{s_8 s_9} \\ m_{22} &= \frac{s_5 s_8}{s_6 s_7} \end{aligned}$$

are all positive Laurent polynomials.

For Λ_2 , we start with the diagram for Λ_1 in Figure 30, and add one more crossing labeled a_{14} directly to the right of a_{13} . Choose the filling L of Λ_2 given by resolving the crossings

$$a_{10}, a_9, a_8, a_7, a_6, a_5, a_{12}, a_{11}$$

in order. The augmentation ε_L sends

$$\begin{aligned} t_1 &\rightarrow -s_6 s_7 s_{10} s_{11}, t_2 \rightarrow -\frac{s_5 s_8 s_9 s_{12}}{s_6 s_7 s_{10} s_{11}}, a_9 \rightarrow s_9, a_{10} \rightarrow s_{10}, a_{13} \rightarrow \frac{s_5}{s_6} - \frac{s_7}{s_8} + \frac{s_9}{s_{10}} - \frac{s_5 s_8 s_9}{s_6^2 s_7^2 s_{10}^2 s_{11}}, \\ a_{14} &\rightarrow \frac{s_6^2 s_7^2 s_{10}^2 s_{11}}{s_5 s_8 s_9} - \frac{s_6^3 s_7^3 s_{10}^3 s_{11}^2}{s_5^2 s_8^3 s_9^3 s_{12}} + \frac{s_6^2 s_7^2 s_{10}^3 s_{11}^2}{s_5 s_8^2 s_9^3 s_{12}} - \frac{s_6^2 s_7^2 s_{10}^3 s_{11}^2}{s_5 s_8 s_9^2 s_{12}}. \end{aligned}$$

Define $M_0 = \varepsilon_L \begin{pmatrix} t_1 & a_{14}t_2 \\ a_{13}t_1 & (1 + a_{13}a_{14})t_2 \end{pmatrix}$, $N = \varepsilon_L \begin{pmatrix} a_{10} & 1 \\ a_9 & 0 \end{pmatrix}$, and $M_1 = N^{-1}M_0N$; then

$M_1|_{\{s_j \rightarrow -s_j, j=5,7,8,10\}} = \begin{pmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{pmatrix}$, and the entries

$$\begin{aligned} m_{11} &= \frac{s_6^2 s_{10}^2 s_{11} s_7^3}{s_5 s_8^3 s_9^2} + \frac{2s_6 s_{10}^2 s_{11} s_7^2}{s_8^2 s_9^2} + \frac{s_6^2 s_{10} s_{11} s_7^2}{s_5 s_8^2 s_9} + \frac{s_6 s_{10} s_{12} s_7^2}{s_8} + \frac{s_5 s_{10}^2 s_{11} s_7}{s_8 s_9^2} + \frac{s_6 s_{10} s_{11} s_7}{s_8 s_9} \\ &\quad + s_6 s_9 s_{12} s_7 + s_5 s_{10} s_{12} s_7 + \frac{s_5}{s_6 s_7 s_9} + \frac{1}{s_8 s_9}, \\ m_{12} &= \frac{s_6 s_{10} s_{11} s_7^2}{s_8 s_9} + s_6 s_{11} s_7 + \frac{s_5 s_{10} s_{11} s_7}{s_9} + \frac{s_5 s_8}{s_6 s_7 s_{10}}, \\ m_{21} &= \frac{s_5 s_7 s_{11} s_{10}^3}{s_8 s_9^2} + \frac{2s_6 s_7^2 s_{11} s_{10}^3}{s_8^2 s_9^2} + \frac{s_6^2 s_7^3 s_{11} s_{10}^3}{s_5 s_8^3 s_9^2} + s_5 s_7 s_{12} s_{10}^2 + \frac{s_6 s_7^2 s_{12} s_{10}^2}{s_8} + \frac{s_5 s_{10}}{s_6 s_7 s_9} + \frac{s_{10}}{s_8 s_9}, \\ m_{22} &= \frac{s_5 s_7 s_{11} s_{10}^2}{s_9} + \frac{s_6 s_7^2 s_{11} s_{10}^2}{s_8 s_9} + \frac{s_5 s_8}{s_6 s_7}. \end{aligned}$$

are all positive Laurent polynomials. □




FIGURE 31. Lagrangian projection for the Legendrian link Λ_n , $n \geq 1$.




FIGURE 32. Lagrangian projection for the Legendrian link $\Lambda(\tilde{D}_n)$, $n \geq 4$, as used in the proof of Theorem 6.2.

6.5. The general case: $\Lambda(\tilde{D}_n)$ and Λ_n . We now turn to the Legendrian links $\Lambda(\tilde{D}_n)$, $n \geq 5$, and Λ_n , $n \geq 3$, which are depicted in Figures 31 and 32. The action of the ϑ -loops on the Legendrian contact DGA for $\Lambda(\tilde{D}_n)$, respectively Λ_n , can be studied directly thanks to our understanding of the ϑ -loops for the Legendrian braids $\Lambda(\tilde{D}_4)$, respectively Λ_2 . The main ingredient that allows us to deduce the general cases from a particular case is the following:

Proposition 6.5. *Let $\beta \in \text{Br}_N^+$ be an admissible braid, let $k_1, k_2 \in \mathbb{N}$ with $2 \leq k_1 \leq k_2$, and let L be an exact Lagrangian filling of $\Lambda(\beta, 1; \sigma_1^{k_1})$. Consider an exact Lagrangian cobordism Σ from $\Lambda(\beta, 1; \sigma_1^{k_1})$ to $\Lambda(\beta, 1; \sigma_1^{k_2})$ obtained by resolving any combination of $(k_2 - k_1)$ crossings in the braid $\sigma_1^{k_2}$ which are not the initial nor the final crossings. Let*

$$\Phi_\Sigma : \mathcal{A}(\Lambda(\beta, 1; \sigma_1^{k_2})) \rightarrow \mathcal{A}(\Lambda(\beta, 1; \sigma_1^{k_1}))$$

denote the induced map between the Legendrian contact DGAs. Let a denote any Reeb chord of $\Lambda(\beta, 1; \sigma_1^{k_1})$ not in $\sigma_1^{k_1}$, as well as the corresponding Reeb chord of $\Lambda(\beta, 1; \sigma_1^{k_2})$. Then for all $m \in \mathbb{N}$, we have

$$(\varepsilon_L \circ \vartheta_1^m)(a) = (\varepsilon_L \circ \Phi_\Sigma \circ \vartheta_2^m)(a),$$

where ϑ_i denotes the ϑ -loop of $\Lambda(\beta, 1; \sigma_1^{k_i})$, $i = 1, 2$.

Consequently, if the ϑ_1 -orbit of the augmentation ε_L of $\Lambda(\beta, 1; \sigma_1^{k_1})$ is entire, then the ϑ_2 -orbit of the augmentation $\varepsilon_{L \circ \Sigma} = \varepsilon_L \circ \Phi_\Sigma$ of $\Lambda(\beta, 1; \sigma_1^{k_2})$ is entire.

Proof. We begin by noting that f_Σ fixes any Reeb chord of $\Lambda(\beta, 1; \sigma_1^{k_2})$ outside of the braid $\sigma_1^{k_2}$. This is because f_Σ consists of a composition of saddle cobordism maps that count disks with two positive corners, and the only such disks with a positive corner at one of the resolved crossings must have its other positive corner at a crossing in the braid, by our assumption that we do not resolve the two extreme crossings of $\sigma_1^{k_2}$.

There are two types of crossings in $\Lambda(\beta, 1; \sigma_1^{k_2})$ besides the crossings in $\sigma_1^{k_2}$: the ones that come from crossings of $\Lambda(\beta)$ involving the satellited strand of β , and the ones that do not. If a is of the latter type, then ϑ_2 and ϑ_1 both fix a . Since $f_\Sigma(a) = a$, we are done in this case.

Now assume that a is of the former type, and note that crossings of this type come in pairs corresponding to the two strands of $\sigma_1^{k_2}$. Recall that the action of the ϑ -monodromy on such a pair of crossings is completely determined by the path matrix of the braid. If we write P_1 and P_2 for the path matrices for the braids $\sigma_1^{k_1}$ and $\sigma_1^{k_2}$ respectively, then it suffices to show that

$$P_1 = f_\Sigma(P_2).$$

Note that these path matrices incorporate all base points in the braid region; in particular, the braid $\sigma_1^{k_1}$ includes base points in its interior, coming from the resolved crossings of $\sigma_1^{k_2}$.



FIGURE 33. Resolving a crossing in σ_1^3 to produce σ_1^2 .

To prove $P_1 = f_\Sigma(P_2)$, by functoriality we may assume that Σ consists of a single saddle cobordism. Furthermore, since f_Σ fixes any crossing besides the two crossings adjacent to the saddle, it suffices to check the equality when $k_1 = 3$ and $k_2 = 2$; see Figure 33. In this case, if a_1, a_2, a_3 denote the crossings in σ_{k_1} as shown in Figure 33, we have $f_\Sigma(a_1) = a_1 - s_2^{-1}$, $f_\Sigma(a_2) = s_2$, $f_\Sigma(a_3) = a_3 - s_2^{-1}$, and we compute:

$$f_\Sigma(P_2) = f_\Sigma \left(\begin{pmatrix} 0 & 1 \\ 1 & a_1 \end{pmatrix} \begin{pmatrix} 0 & 1 \\ 1 & a_2 \end{pmatrix} \begin{pmatrix} 0 & 1 \\ 1 & a_3 \end{pmatrix} \right) = \begin{pmatrix} 0 & 1 \\ 1 & a_1 \end{pmatrix} \begin{pmatrix} -s_2^{-1} & 0 \\ 0 & s_2 \end{pmatrix} \begin{pmatrix} 0 & 1 \\ 1 & a_3 \end{pmatrix} = P_1,$$

as desired.

The final sentence of the proposition follows from the fact that Φ_Σ is surjective; see the proof of Proposition 7.5 below. \square

From Proposition 6.5, and the fact that $\Lambda(\tilde{D}_4)$ and Λ_2 have fillings for which the ϑ -orbit of the associated augmentation is entire, it follows that the same is true for $\Lambda(\tilde{D}_n)$, $n \geq 5$, and Λ_n , $n \geq 3$. This completes the proof of Theorem 1.1.

Remark 6.6. A consequence is that the Legendrian links $\Lambda(\tilde{D}_n)$, $n \geq 5$, and Λ_n , $n \geq 3$ have infinitely many fillings. This conclusion also follows from just the existence of a cobordism from $\Lambda(\tilde{D}_4)$ or Λ_2 to these links, using Proposition 7.5 below. However, Theorem 1.1 is stronger: we have actually constructed an infinite family of fillings of each of these links that are all provably distinct from each other. \square

7. PROOF OF COROLLARIES AND CONCLUDING REMARKS

In this section we discuss some of the applications stated in the Introduction. First, we show that the smooth isotopy type of the Lagrangian fillings we construct is independent of the iteration of the ϑ -loop. Then, we precisely state the notion of *aug-infinite* Legendrians (which implies the existence of infinitely many fillings) and prove some of its properties under exact Lagrangian cobordisms. We also conclude Proposition 7.7, providing a gamut of small smooth knots with a max-tb Legendrian representative that admits infinitely many Lagrangian fillings. Finally, we prove Corollaries 1.6 and 1.7 regarding closed Lagrangian surfaces in certain Weinstein 4-manifolds.

7.1. Smooth isotopy class of Lagrangian fillings. Let $L \subset (\mathbb{D}^4, \lambda_{st})$ be an exact Lagrangian filling of $\Lambda \subset (\mathbb{S}^3, \xi_{st})$ and $\vartheta : \mathbb{S}^1 \rightarrow \mathcal{L}(\Lambda)$ a Legendrian loop. The isotopy cobordism $gr(\vartheta) \subset \mathbb{S}^3 \times [0, 1]$ associated to the Legendrian loop ϑ is an exact Lagrangian self-concordance of Λ , which we can concatenate with L . This yields another exact Lagrangian filling $L_\vartheta = L \# gr(\vartheta)$ of Λ . Theorem 1.1 shows that, for certain ϑ , L_ϑ may *not* be Hamiltonian isotopic to L . For the Legendrian ϑ -loops we use in this article, let us prove that L_ϑ is always *smoothly* isotopic to L . The argument is the same as in (the updated version of) [CG21]; we reproduce it here for convenience:

Proposition 7.1. *Let $\Lambda \subset (\mathbb{S}^3, \xi_{st})$ be a Legendrian link of the form $\Lambda = \Lambda(\beta, i; \gamma)$, where $\beta \in \text{Br}_N^+$, $\gamma \in \text{Br}_M^+$. Let $L_\pi \subset (\mathbb{D}^4, \lambda_{st})$ be an exact Lagrangian filling obtained by a pinching sequence $\pi \in S_{|\beta|}$, and $\vartheta : \mathbb{S}^1 \rightarrow \mathcal{L}(\Lambda)$ a Legendrian ϑ -loop. Then, the exact Lagrangian fillings L and L_ϑ are smoothly isotopic relative to their boundary Λ .*

Proof. From the perspective of a positive braid representative $\beta \in \text{Br}_N^+$ of $\Lambda = \Lambda(\beta)$, a Legendrian ϑ -loop consists of two moves: Reidemeister III moves and conjugations. Let us denote the ordered crossings of β by (a_j) , $j \in [1, |\beta|]$, with $a_{\pi(k)}$ being the k -th crossing to be resolved. First, *any* pinching (resolution) sequence $\pi \in S_{|\beta|}$ yields a surface which is smoothly isotopic to L . From the smooth perspective, resolving a crossing corresponds to an elementary surface cobordism of index 1 and thus two consecutive such cobordisms can be performed in either order without affecting the smooth type, see Figure 34 below. Since any two different pinching sequences differ by a composition of transpositions, the smooth isotopy class of L_π is equal for *any* pinching sequence $\pi \in S_{|\beta|}$. It thus suffices to consider the case of the identity permutation $\pi = e$, and show that $L_e \cup_\Lambda gr(\vartheta)$ is smoothly isotopic to L_e .

Consider a Reidemeister III move for three (consecutive) crossings a_{i-1}, a_i, a_{i+1} , which leads to a_{i+1}, a_i, a_{i-1} . For the Lagrangian filling L_e , these three crossings are resolved left to right: starting at a_{i-1} , then a_i and a_{i+1} , in this order. Starting at a_{i+1}, a_i, a_{i-1} we can describe *two* smooth cobordisms, both local to this piece of the braid (constant relative to its endpoints):




FIGURE 34. A compactly supported smooth isotopy between two (local) exact Lagrangian cobordisms which are not Hamiltonian isotopic. The coordinate $t \in \mathbb{R}$ represents the symplectization direction $\mathbb{R}_t \times R_{x,y,z}^3$ and the diagrams are Lagrangian projections in $\mathbb{R}_{x,y}^2$ as indicated. The variable $s \in [0, 1]$ is the real coordinate associated to the isotopy itself.

- (i) Apply a Reidemeister III move down to a_{i-1}, a_i, a_{i+1} and then resolve according to $\pi = e$. Namely, first a_{i-1} , then a_i and finally a_{i+1} .
- (ii) Directly resolve the three crossings a_{i+1}, a_i, a_{i-1} , using the transposition $\pi = (i, i+2)$. That is, we resolve a_{i+1} first, then a_i and lastly a_{i-1} .

These cobordisms are depicted in Figure 35. Both tangle cobordisms start at the tangle a_{i+1}, a_i, a_{i-1} and end up in the trivial 3-stranded tangle. Since the crossings a_{i-1} and a_{i+1} are interchanged in a Reidemeister III move (a_{i-1} before being geometrically the same as a_{i+1}), as are a_{i+1}, a_{i-1} , the two tangle cobordisms are smoothly isotopic. Hence, concatenating with the graph of an isotopy given by a sequence of Reidemeister III moves, from Λ to itself, does not affect the smooth isotopy type of a Lagrangian filling L .




FIGURE 35. The two (local) exact Lagrangian cobordisms associated to a Reidemeister III move: the cobordism in Item (i) in the text is depicted on the left, whereas the cobordism in Item (ii) is drawn on the right. These cobordisms are to be compared smoothly, relative to their common ends.

The same occurs for conjugation of the given positive braid $\beta = (a_1, a_2, \dots, a_{|\beta|-1}, a_{|\beta}|)$. Indeed, there are two smooth concordances starting with $(a_{|\beta|}, a_1, a_2, \dots, a_{|\beta|-1})$:

- (i) Apply the cyclic shift from $(a_{|\beta|}, a_1, a_2, \dots, a_{|\beta|-1})$ down to $(a_1, a_2, \dots, a_{|\beta|-1}, a_{|\beta}|)$ and then resolve the crossings starting at $a_{|\beta|}$ and then left to right, that is, continuing with a_1, a_2 and resolving through $a_{|\beta|-1}$.
- (ii) Directly resolve the crossings of $(a_{|\beta|}, a_1, a_2, \dots, a_{|\beta|-1})$ according to $\pi = e$: starting at $a_{|\beta|}$, then a_1, a_2 through $a_{|\beta|-1}$.

These two concordances yield smoothly isotopic surfaces. In conclusion, starting with the Lagrangian filling $L_e \subset (\mathbb{D}^4, \lambda_{\text{st}})$, the concatenation $L_\vartheta = L \cup_\Lambda gr(\vartheta)$ yields a Lagrangian filling of the form L_π , for a permutation $\pi \in S_{|\beta|}$. Since L_π and L_e are smoothly isotopic, the required statement follows. \square

7.2. Aug-infinite Legendrian links and cobordisms. Here we describe a method for starting with one Legendrian link known to have infinitely many fillings and producing others. First we need to define a condition that implies having infinitely many fillings and is in turn implied in our examples by the ϑ -orbit being entire.

Suppose that Λ is a Legendrian link with a (connected, orientable, exact Lagrangian) filling L of Maslov number 0. As discussed in Section 3.3, L induces a $(2g + 2m - 2)$ -system of augmentations $\varepsilon_L : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}[s_1^{\pm 1}, \dots, s_{2g+2m-2}^{\pm 1}]$, where g is the genus of L and m is the number of components of Λ . Furthermore, up to equivalence (automorphism of $\mathbb{Z}[s_1^{\pm 1}, \dots, s_{2g+2m-2}^{\pm 1}]$), this system is well-defined, independent of choices, and invariant under Hamiltonian isotopy of L . Here, as in Section 3.1, we have assumed in defining the DGA $(\mathcal{A}_\Lambda, \partial)$ that there is one base point on each component of Λ .

There are $2^{2g+2m-2}$ ring morphisms from $\mathbb{Z}[s_1^{\pm 1}, \dots, s_{2g+2m-2}^{\pm 1}]$ to \mathbb{Z} , each sending each s_i to ± 1 . By composing ε_L with these homomorphisms, we obtain $2^{2g+2m-2}$ augmentations from \mathcal{A}_Λ to \mathbb{Z} . In this way, the filling L of Λ induces finitely many \mathbb{Z} -valued augmentations $(\mathcal{A}_\Lambda, \partial) \rightarrow (\mathbb{Z}, 0)$. Note that this continues to hold even if L is not connected: the augmentations induced by a disconnected filling of Λ necessarily annihilate any Reeb chord of Λ whose endpoints lie on different components of the filling, and each component of the filling induces finitely many augmentations of the sublink of Λ given by the boundary of the component.

Definition 7.2. A Legendrian link Λ is *aug-infinite* if the collection of all \mathbb{Z} -valued augmentations $(\mathcal{A}_\Lambda, \partial) \rightarrow (\mathbb{Z}, 0)$ induced by orientable exact Lagrangian fillings of Λ of Maslov number 0, ranging over all possible such fillings, is infinite.

Note that the aug-infinite condition is independent of the choices made along the way, including spin structure, capping paths and operators, and placement of base points. Adding extra base points also does not affect the condition; cf. the proof of Proposition 7.5 below.

The following is an immediate consequence of the fact that each filling induces finitely many \mathbb{Z} -valued augmentations.

Proposition 7.3. *If Λ is aug-infinite then it has infinitely many fillings.*

In the conclusion of Proposition 7.3, *infinitely many Lagrangian fillings* refers to the fact that there are infinitely many Lagrangian fillings up to Hamiltonian isotopy. A priori, they might not be smoothly isotopic. Nevertheless, as proven in Proposition 7.1, this is the case for the Lagrangian fillings we construct with ϑ -loops.

Next we observe that our arguments from Section 6 actually prove that the Legendrians $\Lambda(\tilde{D}_n)$, Λ_n , and $\Lambda(\tilde{A}_2)$ satisfy this strengthened version of having infinitely many fillings.

Proposition 7.4. *The three classes of Legendrian links $\Lambda(\tilde{D}_n)$ ($n \geq 4$), Λ_n ($n \geq 1$), and $\Lambda(\beta_{11}) = \Lambda(\tilde{A}_2), \Lambda(\beta_{12}), \Lambda(\beta_{21})$ are aug-infinite.*

Now we claim that for particular decomposable Lagrangian cobordisms, if the bottom of the cobordism is aug-infinite, then the top is as well. To be precise, we have the following.

Proposition 7.5. *Let Λ_+ and Λ_- be Legendrian links with rotation number 0, and suppose that Λ_- is an aug-infinite Legendrian link. Suppose that the following two properties hold:*

- *The xy projection $\Pi_{xy}(\Lambda_-)$ is obtained from $\Pi_{xy}(\Lambda_+)$ by a sequence of saddle cobordisms at proper contractible Reeb chords of Λ_+ of degree 0;*
- *All Reeb chords of Λ_- (and thus of Λ_+) are in nonnegative degree.*

Then the Legendrian link Λ_+ is aug-infinite.

Proof. It suffices to consider the case where Λ_+ and Λ_- are related by a single saddle move at a Reeb chord a of Λ_+ . Suppose that Λ_+ has m components, and place a base point on each; these base points trace down to Λ_- . As in Section 3.5, we place a pair of base points on Λ_- coming from the saddle at a . Then both \mathcal{A}_{Λ_+} and \mathcal{A}_{Λ_-} are DGAs over $\mathbb{Z}[t_1^{\pm 1}, \dots, t_m^{\pm 1}, s^{\pm 1}]$, and the cobordism gives a map $\Phi : (\mathcal{A}_{\Lambda_+}, \partial) \rightarrow (\mathcal{A}_{\Lambda_-}, \partial)$.

Any filling of Λ_- produces a system of augmentations for Λ_- as in the discussion in Section 3.5; note that now one component of Λ_- has more than one base point, but the construction from Section 3.5 works just as well in this case. From Remark 3.1, adding each extra base point has the effect on $(\mathcal{A}_{\Lambda_-}, \partial)$ of replacing one generator t of the coefficient ring by two generators t', t'' and setting $t = t't''$ in the differential. It follows that there is a two-to-one correspondence between \mathbb{Z} -valued augmentations of $(\mathcal{A}_{\Lambda_-}, \partial)$ after and before the extra base point is added, and so adding extra base points does not affect the aug-infinite condition. We conclude that $(\mathcal{A}_{\Lambda_-}, \partial)$ has infinitely many augmentations coming from fillings. Since there are finitely many choices for the images of t_1, \dots, t_m, s under such an augmentation, there exist $t_1^0, \dots, t_m^0, s^0 \in \{\pm 1\}$ such that $(\mathcal{A}_{\Lambda_-}, \partial)$ has infinitely many augmentations from fillings that send t_i to t_i^0 and s to s^0 .

Write $\mathcal{A}_{\Lambda_+}^{\mathbb{Z}}$ and $\mathcal{A}_{\Lambda_-}^{\mathbb{Z}}$ for the DGAs over \mathbb{Z} obtained by setting $t_i = t_i^0$ and $s = s^0$. The cobordism map Φ induces a map $\Phi^{\mathbb{Z}} : \mathcal{A}_{\Lambda_+}^{\mathbb{Z}} \rightarrow \mathcal{A}_{\Lambda_-}^{\mathbb{Z}}$ satisfying $\Phi^{\mathbb{Z}}(a) = \pm 1$ and for all other Reeb chords a_i of Λ_+ , $\Phi^{\mathbb{Z}}(a_i) = a_i + f(a_i)$ for some $f(a_i) \in \mathcal{A}_{\Lambda_-}$ determined by the construction in Section 4.2. As observed in Section 4.2, this map respects the height filtration: for each i , $f(a_i)$ only involves Reeb chords of strictly smaller height than a_i . We conclude from this that $\Phi^{\mathbb{Z}}$ is surjective.

Since $\mathcal{A}_{\Lambda_-}^{\mathbb{Z}}$ has infinitely many augmentations from fillings, there is some Reeb chord a_i of Λ_- that is sent to infinitely many values in \mathbb{Z} under these augmentations. Now use the surjectivity of $\Phi^{\mathbb{Z}}$ and suppose that $x \in \mathcal{A}_{\Lambda_+}$ satisfies $\Phi^{\mathbb{Z}}(x) = a_i$. Each augmentation of $\mathcal{A}_{\Lambda_-}^{\mathbb{Z}}$ from a filling of Λ_- produces an augmentation of $\mathcal{A}_{\Lambda_+}^{\mathbb{Z}}$ from a filling of Λ_+ by composition with $\Phi^{\mathbb{Z}}$, and x is sent to infinitely many values in \mathbb{Z} under these augmentations. This shows that $\mathcal{A}_{\Lambda_+}^{\mathbb{Z}}$ has infinitely many augmentations from fillings of Λ_+ , and consequently that Λ_+ is aug-infinite. \square

Remark 7.6. It is expected that Proposition 7.5 should hold whenever there is an exact Lagrangian cobordism between Λ_+ and Λ_- , without the restriction of being composed strictly of saddle moves (and not isotopy cylinders) or even of being decomposable. One approach to proving the more general result is to show that exact cobordisms induce injective maps on the augmentation categories of Legendrian links (over \mathbb{Z}), in the spirit of previous work of Pan [Pan17a] for Legendrian knots and the upcoming paper [CSLL⁺21] for links. \square

As the first application of Proposition 7.5, we have Corollary 1.4:

Proof of Corollary 1.4. Observe that there is a decomposable Lagrangian cobordism to the Legendrian $(4, 4)$ torus link $\Lambda(4, 4)$, which is the (-1) -closure of the 4-braid $(\sigma_1\sigma_2\sigma_3)^8 =$

$(\sigma_2\sigma_1\sigma_3\sigma_2)^4\sigma_3^4\sigma_1^4$, from the link $\Lambda(\tilde{D}_4) = \Lambda((\sigma_2\sigma_1\sigma_3\sigma_2)^4\sigma_3^2\sigma_1^2)$, consisting of two saddle cobordisms at proper contractible degree 0 Reeb chords. Since $\Lambda(\tilde{D}_4)$ is aug-infinite, it follows that $\Lambda(4, 4)$ is aug-infinite as well. In addition, there is another such cobordism from $\Lambda(4, 4)$ to the Legendrian (n, m) torus link $\Lambda(n, m)$ for any $n, m \geq 4$, and so by Proposition 7.5, $\Lambda(n, m)$ is aug-infinite for any $n, m \geq 4$. Similarly, we can deduce that the Legendrian $(3, 6)$ torus link $\Lambda(3, 6)$, which is the (-1) -closure of the 3-braid $(\sigma_1\sigma_2)^9 = (\sigma_2\sigma_1^2\sigma_2)^3\sigma_1^6$, is aug-infinite because there is a cobordism to $\Lambda(3, 6)$ from Λ_1 ; it then also follows that the $(3, m)$ -torus link $\Lambda(3, m)$ is aug-infinite for all $m \geq 6$. The proof is complete. \square

We can also apply Proposition 7.5 to show that various other single-component Legendrian knots have infinitely many fillings.

Proposition 7.7. *The Legendrian knots given by the (-1) -closures of the following positive braids have infinitely many fillings:*

- (i) $(\sigma_2\sigma_1\sigma_3\sigma_2)^4\sigma_2\sigma_1\sigma_3$, which has smooth type $m(10_{145})$, Thurston–Bennequin number 3 and genus 2 fillings,
- (ii) $(\sigma_2\sigma_1\sigma_3\sigma_2)^3\sigma_2^2\sigma_1^2\sigma_3\sigma_2\sigma_1\sigma_3^2 \in B_4$ which has smooth type 10_{154} , Thurston–Bennequin number 5 and genus 3 fillings,
- (iii) $\sigma_2^2\sigma_1^2\sigma_2^2\sigma_1^2\sigma_2^2\sigma_1^2\sigma_2\sigma_1 \in B_3$, which has smooth type $m(10_{161})$, Thurston–Bennequin number 5 and genus 3 fillings,
- (iv) $\sigma_2\sigma_1^2\sigma_2^2\sigma_1^2\sigma_2^2\sigma_1^2\sigma_2^2\sigma_1^3 \in B_3$, of smooth type 10_{139} , Thurston–Bennequin number 7 and genus 4 fillings,
- (v) $\sigma_2\sigma_1^2\sigma_2^2\sigma_1^2\sigma_2^2\sigma_1\sigma_2\sigma_1^2\sigma_2^2\sigma_1 \in B_3$, of smooth type $m(10_{152})$, Thurston–Bennequin number 7 and genus 4 fillings.

Proof. The $m(10_{145})$ and 10_{154} knots have a cobordism from the link $\Lambda(\beta_{11})$, which is the (-1) -closure of $(\sigma_2\sigma_1\sigma_3\sigma_2)^4\sigma_1\sigma_3$. The other three knots have cobordisms from the link Λ_1 , which is the (-1) -closure of $\sigma_2\sigma_1^2\sigma_2^2\sigma_1^2\sigma_2^2\sigma_1^2\sigma_2\sigma_1$. \square

In light of Proposition 7.5, given two Legendrian links Λ_+, Λ_- with infinitely many fillings, we might consider Λ_- to be “simpler” than Λ_+ if there is a saddle cobordism from Λ_- to Λ_+ . Since such a cobordism increases Thurston–Bennequin number as we go from bottom to top, a rough measure of the simplicity of a Legendrian link with infinitely many fillings is given by its Thurston–Bennequin number: the lower the tb , the simpler the link. (Alternatively, we could use $2g + m$ where g is the genus of a connected filling and m is the number of components of the link, since $tb = 2g + m - 2$.) From this perspective, $m(10_{145})$ ($tb = 3$) is the simplest knot that is known to us to have infinitely many fillings, while $\Lambda(\tilde{A}_2)$ ($tb = 2$) is the simplest known link.

Remark 7.8. We presently do not know of any Legendrian knots with infinitely many genus 1 fillings, or of any Legendrian links with infinitely many planar (genus 0) fillings. From the perspective of cluster algebras, the existence of the former would be somewhat unexpected if we restrict to the class of (-1) -closures of admissible braids. \square

7.3. Lagrangian surfaces in Weinstein 4-manifolds. Here we prove Corollaries 1.6 and 1.7. Let $\Lambda \subset (\mathbb{S}^3, \xi_{\text{st}})$ be a Legendrian link with $m := |\pi_0(\Lambda)|$ components, and $W(\Lambda)$ the Weinstein 4-manifold obtained by attaching m Weinstein handles to $(\mathbb{D}^4, \lambda_{\text{st}})$, one along each component of the Legendrian $\Lambda \subset (\mathbb{S}^3, \xi_{\text{st}}) \cong (\partial\mathbb{D}^4, \ker(\lambda_{\text{st}}|_{\partial\mathbb{D}^4}))$. Given an embedded exact Lagrangian filling $L \subset (\mathbb{D}^4, \lambda_{\text{st}})$, we denote by $\bar{L} \subset W(\Lambda)$ the closed embedded exact

Lagrangian surface in $W(\Lambda)$ given by the set-theoretic union $\bar{L} := L \cup L_{cap}$, where L_{cap} is the (disjoint) union of the Lagrangian cores of the m Weinstein handles.

The augmentations $\varepsilon_L : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}[H_1(L)]$ of the Legendrian contact DGA used in this manuscript employ the system of coefficients $\mathbb{Z}[H_1(L)]$, geometrically keeping track of local systems in a Lagrangian filling $L \subset (\mathbb{D}^4, \lambda_{st})$. In the transition from L to $\bar{L} \subset W(\Lambda)$, we must compare $\mathbb{Z}[H_1(L)]$ and $\mathbb{Z}[H_1(\bar{L})]$, which are not isomorphic unless Λ has a single component. This motivates the following definition.

Definition 7.9. Let $L \subset (\mathbb{D}^4, \lambda_{st})$ be a filling of a Legendrian link $\Lambda \subset (\mathbb{S}^3, \xi_{st})$, inducing the system of augmentations $\varepsilon_L : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}[H_1(L)]$, where Λ is equipped with the null-cobordant spin structure. The *restricted system of augmentations* associated to L is the composition

$$\varepsilon_{\bar{L}} : \mathcal{A}_\Lambda \xrightarrow{\varepsilon_L} \mathbb{Z}[H_1(L)] \longrightarrow \mathbb{Z}[H_1(\bar{L})],$$

where the second map is induced by the quotient map $H_1(L) \rightarrow H_1(\bar{L})$. \square

If we place a single base point t_i on each component Λ_i of Λ , then t_i represents the homology class of Λ_i in both $H_1(\Lambda)$ and $H_1(L)$, and the quotient map in Definition 7.9 sends each t_i to 1 since Λ_i is null-homologous in \bar{L} . For practical purposes, if L is a connected decomposable filling of an m -component link Λ , we can compute the restricted system of augmentations $\varepsilon_{\bar{L}}$ associated to L as follows.

Let us write $(\mathcal{A}_\Lambda, \partial)$ for the DGA of Λ with the *Lie group* spin structure, which is a DGA over $\mathbb{Z}[t_1^{\pm 1}, \dots, t_m^{\pm 1}]$. Recall from Sections 3.5 and 3.6 the construction of the system of augmentations $\varepsilon_L : \mathcal{A}_\Lambda \rightarrow R$ where $R = (\mathbb{Z}[t_1^{\pm 1}, \dots, t_m^{\pm 1}, s_1^{\pm 1}, \dots, s_\ell^{\pm 1}]) / (w_1 = \dots = w_k = -1)$. We can further quotient the ring R by the relations $t_1 = \dots = t_m = -1$ to get $\bar{R} := R / (t_1 = \dots = t_m = -1)$: this corresponds to passing from $H_1(L)$ to $H_1(\bar{L})$, where the $-$ sign comes from the fact that we are using the Lie group, rather than the null-cobordant spin structure, on L . From Remark 3.15, $t_1 \cdots t_m = (-1)^m$ in R , and this new quotient imposes $m - 1$ new relations. We have $R \cong \mathbb{Z}[H_1(L) \oplus \mathbb{Z}^{m-1}]$ and $\bar{R} \cong \mathbb{Z}[H_1(\bar{L}) \oplus \mathbb{Z}^{m-1}]$; imposing the conditions $t_1 = \dots = t_m = -1$ on the system of augmentations for L produces the restricted system of augmentations $\varepsilon_{\bar{L}}$ for the Lagrangian filling L , enhanced by link automorphisms in each case.

Remark 7.10. When Λ is a single-component Legendrian knot, there is no difference between the system and the *restricted* system of augmentations for a filling L . This comes from the result of Levenson [Lev16] that any augmentation in this case must necessarily send the unique t variable to -1 ; geometrically, this correlates with the fact that Λ is already null-homologous in L before we pass to \bar{L} . \square

The purpose of restricted systems of augmentations for L is that they correspond to local systems that extend to local systems for the closed exact Lagrangian surface \bar{L} . Let $L_1, L_2 \subset (\mathbb{D}^4, \lambda_{st})$ be Lagrangian fillings of a Legendrian link $\Lambda \subset (\mathbb{S}^3, \xi_{st})$. If L_1, L_2 are Hamiltonian isotopic, then their associated augmentations are DGA homotopic; see Theorem 3.6. As noted in Remark 3.7, for the Legendrian links studied in this paper, we can replace ‘‘DGA homotopic’’ by a simpler notion. Following Definition 3.9, we define two restricted systems of augmentations

$$\varepsilon_{\bar{L}_1} : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}[H_1(\bar{L}_1)], \quad \varepsilon_{\bar{L}_2} : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}[H_1(\bar{L}_2)]$$

to be equivalent if there exists an isomorphism $\psi : \mathbb{Z}[H_1(\bar{L}_1)] \rightarrow \mathbb{Z}[H_1(\bar{L}_2)]$ such that

$$\varepsilon_{\bar{L}_2} = \psi \circ \varepsilon_{\bar{L}_1}.$$

Then for Legendrian links such that the entire DGA \mathcal{A}_Λ is concentrated in nonnegative degree, as is the case for all of the examples in this paper, DGA homotopic (restricted) systems of augmentations are necessarily equivalent (restricted) systems of augmentations.

Now, both Corollaries 1.6 and 1.7 will be proven by using the following Proposition 7.11, which is not essentially new and uses the recent articles [EL19, Ekh19, GPS19]. We thank T. Ekhholm, S. Ganatra, and Y. Lekili for illuminating discussions regarding the proof of Proposition 7.11.

Proposition 7.11. *Let $\Lambda \subset (\mathbb{S}^3, \xi_{st})$ be a Legendrian link and $L_1, L_2 \subset (\mathbb{D}^4, \lambda_{st})$ two fillings of Λ . Suppose that the two restricted systems of augmentations*

$$\varepsilon_{\bar{L}_1} : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}[H_1(\bar{L}_1)], \quad \varepsilon_{\bar{L}_2} : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}[H_1(\bar{L}_2)]$$

are not DGA homotopic. Then the exact Lagrangian surfaces $\bar{L}_1, \bar{L}_2 \subset W(\Lambda)$ are not Hamiltonian isotopic in the Weinstein 4-manifold $W(\Lambda)$.

Proof. Let $\mathfrak{W}(W(\Lambda))$ be the wrapped Fukaya category of the Weinstein 4-manifold $W(\Lambda)$, C the union of the m co-cores of the Weinstein handles of $W(\Lambda)$, and $\text{CW}(C)$ the endomorphism ring of C as an object in $\mathfrak{W}(W(\Lambda))$. By [CDGG17, Theorem 1.1], C generates $\mathfrak{W}(W(\Lambda))$, see also [GPS19, Theorem 1.10], and thus we consider the category $\mathfrak{W}(W(\Lambda))$ through its Yoneda embedding $\text{Hom}(C, -) := \text{CW}(C, -)$. The Lagrangian surfaces \bar{L}_1, \bar{L}_2 are exact and hence represent objects in $\mathfrak{W}(W(\Lambda))$, equally denoted \bar{L}_1, \bar{L}_2 . Under the Yoneda embedding, these two objects become $\text{Hom}(C, \bar{L}_i) := \text{CW}(C, \bar{L}_i)$, $i = 1, 2$. We will now argue that $\bar{L}_1, \bar{L}_2 \in \text{Ob}(\mathfrak{W}(W(\Lambda)))$ are distinct objects, which proves that the exact Lagrangian surfaces $\bar{L}_1, \bar{L}_2 \subset W(\Lambda)$ are not Hamiltonian isotopic. It suffices to show that $\text{CW}(C, \bar{L}_1)$ and $\text{CW}(C, \bar{L}_2)$ are distinct as $\text{CW}(C)$ -modules.

Let $L \subset (\mathbb{D}^4, \lambda_{st})$ be a filling of Λ and $\varepsilon_L : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}[H_1(L)]$ its associated augmentation. The holomorphic disks that define ε_L are explained in detail in [EHK16], see also [EN18, Theorem 6.8] and Sections 3 and 4 above. In short, a Reeb chord $a \in \mathcal{A}_\Lambda$ is sent to the contributions from rigid holomorphic disks $u : (\mathbb{D}^2, \partial\mathbb{D}^2) \rightarrow (\mathbb{D}^4, \lambda_{st})$ with a positive puncture at the Reeb chord a , and each disk contribution is weighted by the homology class $[\partial u] \in \mathbb{Z}[H_1(L)]$, where $\partial u \subset L$ is appropriately capped in L . The claim is that the holomorphic disks that define the $\text{CW}(C)$ -module structure of $\text{CW}(C, \bar{L})$, namely the composition A_∞ -map

$$\eta_L : \text{CW}(C) \otimes \text{CW}(C, \bar{L}) \rightarrow \text{CW}(C, \bar{L}),$$

or equivalently $\text{CW}(C) \rightarrow \text{End}(\text{CW}(C, \bar{L}))$, are in bijection with those contributing to the restricted augmentation ε_L . Indeed, we first observe that $C \cap \bar{L}$, which generates $\text{CW}(C, \bar{L})$, consists of precisely a point per each component of C . The disks contributing to η_L have: a positive puncture at a generator of $\text{CW}(C)$, which is either a minimum of a Morse function on C or a Reeb chord of its Legendrian boundary $\partial C \subset \partial W(\Lambda)$; a positive puncture at a generator of $\text{CW}(C, \bar{L})$; and a negative puncture at a generator of $\text{CW}(C, \bar{L})$ (in fact, the two generators of $\text{CW}(C, \bar{L})$ here must be the same). These disks are depicted in the right diagram of Figure 36. In our case, the contributions of these disks are weighted by their boundary homology classes, where we only keep track of the piece of the boundary that belongs to the closed Lagrangian surface \bar{L} . These contributions yield coefficients in the ground ring $\mathbb{Z}[H_1(\bar{L})]$.

Now, the decomposition $\bar{L} = L \cup_\Lambda L_{cap}$ of the Lagrangian surface $\bar{L} \subset W(\Lambda)$ into a Lagrangian filling L and the cores L_{cap} is compatible with neck-stretching along the contact hypersurface $(\partial\mathbb{D}^4, \xi_{st})$ containing Λ , where the Weinstein handles are attached. That is, the Weinstein 4-manifold decomposes as

$$W(\Lambda) \cong (\mathbb{D}^4, \lambda_{st}) \cup_{\mathcal{O}_P(\Lambda)} \left(\bigcup_{i=1}^l (T^*\mathbb{D}^2, \lambda_{st}) \right),$$

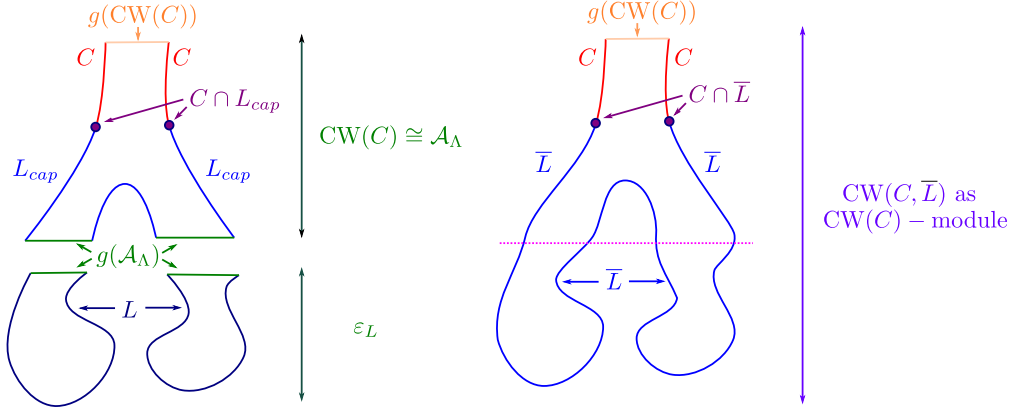


FIGURE 36. On the left, the holomorphic disks contributing to the augmentation ε_L (bottom) and the surgery isomorphism $\text{CW}(C) \cong \mathcal{A}_\Lambda$ (top). On the right, the holomorphic disks contributing to the module structure $\text{CW}(C) \rightarrow \text{End}(\text{CW}(C, \bar{L}))$. The notation $g(\text{CW}(C))$ and $g(\mathcal{A}_\Lambda)$ stands for generators of the algebras $\text{CW}(C)$ and \mathcal{A}_Λ : $g(\text{CW}(C))$ are Reeb chords of ∂C or the minimum of C , and $g(\mathcal{A}_\Lambda)$ are Reeb chords of Λ .

and performing a neck-stretching procedure to the holomorphic disks contributing to η_L breaks them into two pieces. See [Abb14, Chapter 3], [BEH⁺03, Section 3], or [CDGG20, Section 5] for the neck-stretching technique along such a contact hypersurface, in this case a standard contact level set of the symplectization of $(\mathbb{S}^3, \xi_{\text{st}})$. For a sufficiently large stretching, see e.g. [EHK16, Corollary 3.10] or [BEH⁺03, Section 11.3], there is a one-to-one correspondence between the rigid holomorphic disks contributing to η_L and two-level broken disks.

The first level consists of holomorphic disks in the moduli space $\mathcal{M}^{\text{co}}(\mathbf{c})$, following the notation in [EL19], where $\mathbf{c} := c_\Lambda z^v c z^w$, c_Λ is a product of Reeb chords in Λ , z^v, z^w are intersections in $C \cap L_{\text{cap}}$, and c is a generator of $\text{CW}(C)$. The boundaries of these holomorphic disks start belonging in L_{cap} , at the left of the leftmost positive puncture in c_Λ , then continue to belong to L_{cap} as the Reeb chords in c_Λ are visited, and switch to belonging to C , when z^v is reached; then, the boundary (away from the punctures) belongs to C as the chords in c are visited and we reach z^w , where the boundary switches back to L_{cap} . The curves in this first level are depicted in the top of the left diagram of 36, where z^v, z^w are the two points marked by $C \cap L_{\text{cap}}$, and these moduli were studied in detail in [EL19] by using the properties proved in [Ekh19]. In particular, [EL19, Theorem 2] shows that the A_∞ -map $\{\Phi_i\}_{i \in \mathbb{N}} : \text{CW}(C)^{\otimes i} \rightarrow \mathcal{A}_\Lambda$ defined by counting rigid contributions of the moduli spaces $\mathcal{M}^{\text{co}}(\mathbf{c})$ (for $i = 1$; for $i > 1$, we have multiple positive punctures at generators of $\text{CW}(C)$) is an A_∞ -quasi-isomorphism, see [EL19, Theorem 72] for details.

The second level consists of holomorphic disks with a positive puncture at the Reeb chords of Λ and boundary in L . These are the same rigid holomorphic disks as those contributing to the augmentation map $\varepsilon_L : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}[H_1(L)]$. However, we note that the weights are counted with coefficients in $\mathbb{Z}[H_1(\bar{L})]$; that is, the count of holomorphic disks contributing to the second level of η_L is precisely given by the restricted augmentation $\varepsilon_{\bar{L}} : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}[H_1(\bar{L})]$.

In conclusion, the moduli space of disks contributing to the $\text{CW}(C)$ -module structure η_L splits into $\mathcal{M}^{\text{co}}(\mathbf{c})$, which yields the A_∞ -quasi-isomorphism $\text{CW}(C) \cong \mathcal{A}_\Lambda$, and the moduli space of holomorphic disks contributing to the restricted augmentation $\varepsilon_{\bar{L}}$, associated to the closed Lagrangian \bar{L} . Thus, under the surgery isomorphism, the $\text{CW}(C)$ -module structure η_L is precisely given by the augmentation $\varepsilon_{\bar{L}}$ on \mathcal{A}_Λ .

Now if L_1, L_2 are Hamiltonian isotopic fillings of Λ , then \bar{L}_1, \bar{L}_2 are (exact) Lagrangian isotopic in $W(\Lambda)$, even relative to the co-cores. It follows that their associated restricted augmentations $\varepsilon_{\bar{L}_1}$ and $\varepsilon_{\bar{L}_2}$ are DGA homotopic, and the result follows. \square

For the following two corollaries, we emphasize, and implicitly use, that the DGA \mathcal{A}_Λ associated to the Legendrian braids $\Lambda = \Lambda(\beta) \subset (\mathbb{S}^3, \xi_{\text{st}})$, $\beta \in \text{Br}_N^+$, are concentrated in nonnegative degree, and thus DGA homotopic (restricted systems of) augmentations are the same as equivalent (restricted systems of) augmentations.

Proof of Corollary 1.6. For $g = 2$, we consider the Legendrian knot $\Lambda = \Lambda(\beta) \subset (\mathbb{S}^3, \xi_{\text{st}})$ given by the positive braid $\beta = (\sigma_2\sigma_1\sigma_3\sigma_2)^4\sigma_2\sigma_1\sigma_3$. By Proposition 7.7, $\Lambda(\beta)$ admits infinitely many genus 2 exact Lagrangian fillings $\{L_i\}_{i \in \mathbb{N}}$, distinguished by their augmentations $\varepsilon_{L_i} : \mathcal{A}_\Lambda \rightarrow \mathbb{Z}[H_1(L_i)]$. Consider the Weinstein 4-manifold $W := W(\Lambda)$, which is homotopic to a 2-sphere \mathbb{S}^2 because $\Lambda(\beta)$ is a knot. For the same reason, all Lagrangian fillings of Λ are restricted. Note that since Λ is a knot, the restricted augmentation $\varepsilon_{\bar{L}_i}$ is the same as ε_{L_i} for all i . By Proposition 7.11, it follows that the exact Lagrangian surfaces $\{\bar{L}_i\}_{i \in \mathbb{N}}$ in W are not Hamiltonian isotopic. This proves the assertion in the case of $g = 2$.

For higher $g \geq 2$, it suffices to apply the same argument to the Legendrian knots associated to the braids $\beta_g = (\sigma_2\sigma_1\sigma_3\sigma_2)^4\sigma_2\sigma_1\sigma_3\sigma_1^{2(g-2)}$. Since there exists an exact Lagrangian cobordism from $\Lambda(\beta) = \Lambda(\beta_2)$ to $\Lambda(\beta_g)$ for all $g \geq 2$, each knot $\Lambda(\beta_g)$ admits infinitely many exact Lagrangian fillings of genus g . Hence Proposition 7.11 implies that the Weinstein 4-manifold $W_g := W(\Lambda(\beta_g))$, homotopic to a 2-sphere \mathbb{S}^2 , also admits infinitely many exact Lagrangian surfaces of genus g which are not Hamiltonian isotopic. In each case, W_g does not admit any embedded exact Lagrangian surface of genus $h \leq g - 1$ since its intersection form is given by the 1×1 matrix $(\text{tb}(\Lambda(\beta_g)) - 1)$. This concludes the proof. \square

Proof of Corollary 1.7. Consider the Legendrian link $\Lambda = \Lambda(\beta_{11})$ and the Weinstein 4-manifold $W(\Lambda(\beta_{11}))$, which is homotopic to $\mathbb{S}^2 \vee \mathbb{S}^2$ because $\Lambda(\beta_{11})$ has two components. Theorem 1.1 implies that this 2-component link admits infinitely many distinct exact Lagrangian fillings. In order to apply Proposition 7.11, we need to ensure that these infinitely many fillings are distinguished by their restricted systems of augmentations. For that, let us study the augmentation ε_L associated to the (initial) Lagrangian filling L in Subsection 6.3 and its ϑ -loop iterates. There are four homology variables t_1, t_2, t_3, t_4 ; under ε_L , these are augmented to

$$t_1 \rightarrow \frac{s_9 s_{12} s_{13}}{s_{11}}, \quad t_2 \rightarrow -s_{11} s_{16}, \quad t_3 \rightarrow -\frac{1}{s_{10} s_{16}}, \quad t_4 \rightarrow \frac{s_{10}}{s_9 s_{12} s_{13}}.$$

Note that the ϑ -loop monodromy fixes each homology variable, and so the ϑ -loop iterates $\varepsilon_L \circ \vartheta^k$ have the same effect as ε_L on t_1, t_2, t_3, t_4 for all $k \in \mathbb{N}$.

The first two variables t_1, t_2 lie in one component of Λ and t_3, t_4 lie in the other component. From the discussion following Definition 7.9, we can impose the additional conditions $(t_1, t_2, t_3, t_4) = (+1, -1, +1, -1)$ to obtain the restricted system of augmentations $\varepsilon_{\bar{L}}$ (enhanced by link automorphisms); this is because we first set $t_1 = t_3 = 1$ to reduce to a single base point on each component, and then set $t_2 = t_4 = -1$ to pass from $H_1(L)$ to $H_1(\bar{L})$. In terms of the s variables, there are 3 new conditions (the 4th is redundant):

$$s_9 s_{12} s_{13} = s_{11}, \quad s_{11} s_{16} = 1, \quad s_{10} s_{16} = -1.$$

Now we note that $(s_9, s_{10}, s_{11}, s_{12}, s_{13}, s_{16}) = (1, 1, -1, -1, 1, -1)$ in particular satisfy these conditions. These values of the s_i also produce the maximal value of $|(\varepsilon_L \circ \vartheta^k)(a_9)|$ for all $k \in \mathbb{N}$, from the computation in Section 6.3. It follows that the same argument that we used there, to show that the ϑ -orbit of the system of augmentations ε_L is entire, also shows that the same is true of the restricted system of augmentations $\varepsilon_{\bar{L}}$. We can now apply Proposition 7.11 to conclude that $W(\Lambda(\beta_{11}))$ admits infinitely many distinct exact Lagrangian tori, up to Hamiltonian isotopy. \square

APPENDIX A. THE COBORDISM MAP FOR AN ELEMENTARY SADDLE COBORDISM

The goal of this section is to prove Proposition 4.8, the formula for the cobordism map over \mathbb{Z} for a saddle cobordism at a proper contractible Reeb chord. The proof will come in several steps. In Section A.1, we will first add what we call “mini-dips” on either side of the Reeb chord, which then propagate through the cobordism; this changes the cobordism by a Hamiltonian isotopy. The advantage of adding these mini-dips is that they localize the disks that contribute to the cobordism map, so that the map mod 2 is quite simple and can be written down very explicitly. The main technical result is lifting this map to \mathbb{Z} and showing that Proposition 4.8 holds for the cobordism with mini-dips; this is the content of Proposition A.1 below. The proof of Proposition A.1 is somewhat indirect and involves making the cobordism even more complicated, with the trade-off benefit being that the cobordism map becomes easier to handle. This is in the spirit of a well-known technique in Legendrian knot theory called “dipping”, and occupies Sections A.2 and A.3. Finally, in Section A.4, we deduce Proposition 4.8 from its mini-dipped special case, Proposition A.1, by tracing the effect of mini-dips on the cobordism map.

A.1. Formula for the saddle cobordism map.




FIGURE 37. A saddle cobordism (left) and the mini-dipped version of this cobordism (right).

As in Section 4.2, we consider a saddle cobordism L_a between Legendrian links Λ_- and Λ_+ , where Λ_- is obtained from Λ_+ by replacing a contractible Reeb chord a of Λ_+ by the oriented resolution of the crossing. To simplify the cobordism map, we perturb Λ_{\pm} by a Legendrian isotopy (and consequently the saddle cobordism by a Hamiltonian isotopy) as follows: use two Reidemeister II moves to push the understrand of a over the overstrand on either side of a , as shown in Figure 37. We call these moves “mini-dips” of a .¹⁸ Note that the crossing a is situated differently in Figure 37 than the similar-looking Figure 18 from [EHK16], and consequently our mini-dip is different from the dip considered there. Also note that the crossing data for the mini-dips (with the understrand of a passing over the overstrand in the minidips) is forced by the condition that we want the resulting diagrams to represent Lagrangian projections of Legendrian links—apply Stokes’ Theorem to a bigon whose two corners are the contractible chord a and an adjacent crossing in either of the mini-dips.

For the next few subsections, we will assume that Λ_+ and Λ_- contain the mini-dips shown in Figure 37; we will return to the general case without mini-dips in Section A.4. Over \mathbb{Z}_2 ,

¹⁸These are independently introduced in [GSW20a], where they are called “double dipping” and are used for the same purpose of simplifying cobordism maps.

the mini-dips force the cobordism map $\Phi_{L_a} : \mathcal{A}_{\Lambda_+} \rightarrow \mathcal{A}_{\Lambda_-}$ to have the following simple form:

$$\begin{aligned}\Phi_{L_a}(a) &= s \\ \Phi_{L_a}(a_1) &= a_1 + s^{-1} \\ \Phi_{L_a}(a_2) &= a_2 + s^{-1} \\ \Phi_{L_a}(a_i) &= a_i\end{aligned}$$

where the final equation holds for all Reeb chords a_i of Λ_+ besides a, a_1, a_2 . This follows directly from the work of [EHK16] (cf. Section 4.2) because there are only two disks with a positive puncture at a , the bigon between a_1 and a and the bigon between a and a_2 .

Over the course of Sections A.2 and A.3, we will prove that the cobordism map Φ_{L_a} with signs is given as follows.




FIGURE 38. A saddle cobordism with mini-dips. Crossings and base points are labeled, and quadrants with negative orientation sign are shaded.

Proposition A.1. *Suppose that Λ_+ and Λ_- are related by a saddle cobordism as in Figure 38: the Lagrangian projection of Λ_+ has a contractible crossing a flanked by mini-dips with the crossings on either side of a labeled by a_1 and a_2 , and Λ_- is the result of resolving the crossing a and placing base points labeled s and $-s^{-1}$ on either side of the resolved crossing. Then, the cobordism map $\Phi_{L_a} : \mathcal{A}_{\Lambda_+} \rightarrow \mathcal{A}_{\Lambda_-}$ over \mathbb{Z} is given, up to a link automorphism of Λ_- , by:*

$$\begin{aligned}\Phi(a) &= s \\ \Phi(a_1) &= a_1 - s^{-1} \\ \Phi(a_2) &= a_2 - s^{-1} \\ \Phi(a_i) &= a_i\end{aligned}$$

where the final equation holds for all Reeb chords a_i of Λ_+ besides a, a_1, a_2 . More precisely, there is a link automorphism $\Omega : \mathcal{A}_{\Lambda_-} \rightarrow \mathcal{A}_{\Lambda_-}$ such that the $\Phi_{L_a} : \mathcal{A}_{\Lambda_+} \rightarrow \mathcal{A}_{\Lambda_-}$ is chain homotopy equivalent to $\Omega \circ \Phi$ with Φ as defined above.

Remark A.2. We believe that the auxiliary data needed to define signs (capping operators, etc.) can be chosen so that the combinatorial formula for Φ in Proposition A.1 is precisely the geometric map Φ_{L_a} , without composing with a link automorphism of Λ_- . However, we will not need the stronger statement for our purposes. \square

A.2. Splashes and diagonal automorphisms. Our strategy for proving Proposition A.1 is as follows: the signs for the formula for Φ_{L_a} given there are essentially forced, up to a link automorphism of Λ_- , by the algebraic requirement that Φ_{L_a} needs to be a chain map over \mathbb{Z} . This forcing is not true in full generality, but we will see that it is true if we isotop Λ_{\pm} via Reidemeister II moves so that their differentials consist of many terms, each of which is easy to handle. This sort of strategy is familiar in the subject through the technique of dipping; see, e.g., [FR11, Sab05, Siv11]. We will present a variant of this technique in this subsection, and then return to the proof of Proposition A.1 in Section A.3 below.

Let Λ be a Legendrian link. By applying planar isotopy and Reidemeister II moves, we can isotop Λ so that its Lagrangian projection $\Pi_{xy}(\Lambda)$ satisfies the following properties:

- all vertical tangencies (parallel to the y axis) lie on two lines $x = c_0$ and $x = c_1$, and there are at least 2 vertical tangencies on each of these lines;
- no crossings in either the Lagrangian or front projections occur at the same x coordinate.

Note in particular that $\Pi_{xy}(\Lambda)$ is the plat closure of some braid between $x = c_0$ and $x = c_1$ where the braid strands go from left to right. Furthermore, in the front projection $\Pi_{xz}(\Lambda)$, all left cusps lie on the line $x = c_0$ and all right cusps lie on $x = c_1$.

Subdivide the interval $[c_0, c_1]$ by choosing $x_0 < x_1 < \dots < x_p$ with $x_0 = c_0$, $x_p = c_1$ such that:

- there are no crossings in either $\Pi_{xy}(\Lambda)$ or $\Pi_{xz}(\Lambda)$ in the intervals $x \in [x_0, x_1]$ and $[x_{p-1}, x_p]$;
- for $i = 1, \dots, p-2$, in the interval $x \in [x_i, x_{i+1}]$ there is exactly one crossing in either $\Pi_{xy}(\Lambda)$ or $\Pi_{xz}(\Lambda)$, and no crossing in the other.




FIGURE 39. A set of splashes, in the front projection (left) and corresponding Lagrangian projection (right).

Now in a neighborhood of the $x = x_i$ slices for $i = 1, \dots, p-1$, introduce a collection of “splashes”¹⁹ as shown in Figure 39. This is a C^0 -small perturbation in the front projection, while in the Lagrangian projection, each strand is pushed through the other strands. For definiteness, we order the collection of splashes at $x = x_i$ from left to right in increasing order of the y -coordinate of the splashed strand; in the Lagrangian projection, the crossing information for the new crossings is determined by the relative z coordinates of the strands at $x = x_i$. Let Λ' denote the resulting Legendrian link, and note that $\Pi_{xy}(\Lambda')$ is obtained from $\Pi_{xy}(\Lambda)$ by a (large) number of Reidemeister II moves. See Figure 40 for a sample illustration of Λ' .

Write the Chekanov–Eliashberg DGA of Λ' as $(\mathcal{A}_{\Lambda'}, \partial)$. Say that an automorphism Ψ of the algebra $\mathcal{A}_{\Lambda'}$ is *diagonal* if it is of the following form: if a_i denote the Reeb chords of Λ' , then there is a collection of (invertible) scalars λ_i such that $\Psi(a_i) = \lambda_i a_i$ for all i .

Proposition A.3. *Suppose that Ψ is a diagonal automorphism of $\mathcal{A}_{\Lambda'}$ that is also a chain map: $\Psi \circ \partial = \partial \circ \Psi$. Then Ψ is a link automorphism of Λ' .*

In order to prove Proposition A.3, we need some more notation. Let s denote the number of vertical tangencies in the Lagrangian projection of Λ at each of $x = c_0$ and $x = c_1$, so that the Lagrangian projection is the plat closure of a $2s$ -stranded braid. Number these strands $1, 2, \dots, 2s$ so that in $[x_0, x_1]$, the strands are numbered in increasing order of y coordinate;

¹⁹This terminology is inspired by [FR11], though our splashes are slightly different from theirs and more resemble what [EHK16] call “dips”.




FIGURE 40. A full example of splashing. Top to bottom: the front and Lagrangian projections of Λ , with strands joining $x = c_0$ and $x = c_1$ labeled 1, 2, 3, 4, and the Lagrangian projection of Λ' with some crossings labeled.

keep the numbering of braid strands consistent throughout the braid, and that in general the strands will not remain numbered in increasing order beyond $x = x_1$.

Now suppose that Ψ satisfies the hypotheses of Proposition A.3. We will construct units u_1, \dots, u_{2s} such that the following condition holds for all Reeb chords a of Λ' :

$$(*) \quad \Psi(a) = u_{r(a)} u_{c(a)}^{-1} a.$$

Here we use $r(a)$ and $c(a)$ to denote the labels of the strands that are the endpoint and beginning point of a , respectively.

The following lemma is a useful tool for propagating condition (*). Say that an embedded bigon with boundary on $\Pi_{xy}(\Lambda')$ and two convex corners at Reeb chords of Λ' is a *standard bigon* if one corner is $+$ and one is $-$; similarly say that an embedded triangle is a *standard triangle* if one corner is $+$ and the other two are $-$.

Lemma A.4. *If a_1, a_2 are Reeb chords of Λ' such that there is a (unique) standard bigon with corners at a_1, a_2 , then (*) holds for a_1 if and only if it holds for a_2 . If a_1, a_2, a_3 are Reeb chords such that there is a (unique) standard triangle with corners at a_1, a_2, a_3 , then if (*) holds for two of a_1, a_2, a_3 , then it holds for the third as well.*

Proof. A bigon with $+$ corner at a_1 and $-$ corner at a_2 contributes a term a_2 to $\partial(a_1)$; since $r(a_1) = r(a_2)$ and $c(a_1) = c(a_2)$ and $\Psi\partial = \partial\Psi$, it follows that if (*) holds for one of a_1, a_2 , then it holds for the other. Similarly, a triangle with $+$ corner at a_1 and $-$ corners at a_2, a_3 contributes a term $a_2 a_3$ to $\partial(a_1)$; now use the fact that $r(a_1) = r(a_2)$, $c(a_2) = r(a_3)$, and $c(a_1) = c(a_3)$ to conclude the desired result. See Figure 41. \square




FIGURE 41. Left two diagrams: a standard bigon and a standard triangle. Right diagram: a chain of bigons joining $a_{ij}^{q,1}, a_{ij}^{q,2}, a_{ji}^{q,1}, a_{ji}^{q,2}$. (If strand i is instead above strand j , then these are still standard bigons but all $+$ and $-$ labels are interchanged.)

We next label the crossings of Λ' in the splashes as follows. Consider the splashed portion of strand i at $x = x_q$. For any $j \neq i$, this splash crosses strand j twice; label these two crossings $a_{ij}^{q,1}$ (left) and $a_{ij}^{q,2}$ (right). In this way we label all splashed crossings in Λ' as $a_{ij}^{q,l}$ for $1 \leq i, j \leq 2s$ ($i \neq j$), $1 \leq q \leq p-1$, and $1 \leq l \leq 2$. See Figure 40 for an example.

Lemma A.5. *For fixed i, j, q , if one of the crossings $a_{ij}^{q,1}, a_{ij}^{q,2}, a_{ji}^{q,1}, a_{ji}^{q,2}$ satisfies $(*)$, then so do the other three.*

Proof. In a neighborhood of $x = x_q$, strand i lies either completely above or completely below strand j in the z coordinate. It follows that there is a chain of three standard bigons linking $a_{ij}^{q,1}, a_{ij}^{q,2}, a_{ji}^{q,1}, a_{ji}^{q,2}$; see Figure 41. The result follows from Lemma A.4. \square

Lemma A.6. *If Ψ satisfies the hypotheses of Proposition A.3, then there are u_1, \dots, u_{2s} such that $(*)$ holds for all Reeb chords of Λ' .*

Proof. We will prove that $(*)$ holds for all $a = a_{ij}^{q,l}$ by induction on q . In the course of the proof, we will also show that $(*)$ holds for all other Reeb chords of Λ' , which correspond precisely to the Reeb chords of Λ .

We first establish the induction base case $q = 1$. Set $u_1 = 1$. Then for $j = 2, \dots, 2s$, the Reeb chords $a_{1j}^{1,1}$ have one endpoint on strand 1 and one endpoint on strand j ; since each $\Psi(a_{1j})^{1,1}$ is an invertible scalar multiple of $a_{1j}^{1,1}$, it follows that there are unique choices of u_2, \dots, u_{2s} so that $(*)$ holds for $a = a_{1j}^{1,1}$ for all $j = 2, \dots, 2s$. Thus by Lemma A.5, $(*)$ also holds for $a_{1j}^{1,l}$ and $a_{j1}^{1,l}$, $j = 2, \dots, 2s$, $l = 1, 2$. Next suppose $j > i \geq 2$. Consider the two triangles shown in Figure 42. Of the two corners at $a_{ij}^{1,2}$, one must be $+$ and one must be $-$, and similarly for the two corners at $a_{i1}^{1,1}$. Of the corner at $a_{ij}^{1,2}$ and the corner at $a_{j1}^{1,1}$, again one must be $+$ and one must be $-$ since the union of the two triangles is a standard bigon. Since no triangle can have three $-$ corners by Stokes' Theorem, it follows that one of the two triangles in Figure 42 must be standard. Thus by Lemma A.4, $(*)$ holds for $a_{ij}^{1,2}$, whence it holds for $a_{ij}^{1,l}$ and $a_{ji}^{1,l}$ by Lemma A.5. This completes the base case $q = 1$.




FIGURE 42. Showing that $a_{ij}^{1,2}$ satisfies $(*)$.

Now suppose that $(*)$ holds for $a = a_{ij}^{q,l}$ for fixed q and all i, j, l ; we need to show that it also holds for $a = a_{ij}^{q+1,l}$ for all i, j, l . There are two cases depending on whether the crossing between $x = x_q$ and $x = x_{q+1}$ is in $\Pi_{xy}(\Lambda)$ or in $\Pi_{xz}(\Lambda)$. First suppose that the crossing is in $\Pi_{xy}(\Lambda)$, and let k_1, k_2 denote the labels of the strands involved in the crossing. Choose any two indices $i \neq j$ and assume without loss of generality that $a_{ji}^{q,2}$ is to the right of $a_{ij}^{q,2}$. As long as $\{i, j\} \neq \{k_1, k_2\}$, there is a standard bigon joining $a_{ji}^{q,2}$ to $a_{ij}^{q+1,1}$, and it follows from Lemmas A.4 and A.5 and the induction hypothesis that $a_{ij}^{q+1,1}, a_{ij}^{q+1,2}, a_{ji}^{q+1,1}, a_{ji}^{q+1,2}$ satisfy $(*)$. If on the other hand $\{i, j\} = \{k_1, k_2\}$, then if we label the crossing between $x = x_q$ and $x = x_{q+1}$ by a , there are standard bigons joining $a_{ji}^{q,2}$ to a and a to $a_{ij}^{q+1,1}$, and it follows as before that $a_{ij}^{q+1,1}, a_{ij}^{q+1,2}, a_{ji}^{q+1,1}, a_{ji}^{q+1,2}$, along with a itself, all satisfy $(*)$.

It remains to treat the case where the crossing between $x = x_q$ and $x = x_{q+1}$ is in $\Pi_{xz}(\Lambda)$. Say that this crossing is between strands k_1 and k_2 , where we choose the labels so that strand k_2 has larger y coordinate than strand k_1 between $x = x_q$ and $x = x_{q+1}$. The only difference between the splashes at $x = x_q$ and $x = x_{q+1}$ is that strand k_1 lies above k_2 at x_q while k_2 lies above k_1 at x_{q+1} , or vice versa. It follows that for any two indices $i \neq j$, as long as $\{i, j\} \neq \{k_1, k_2\}$, there is a standard bigon joining $a_{ji}^{q,2}$ to $a_{ij}^{q+1,1}$ (or $a_{ij}^{q,2}$ to $a_{ji}^{q+1,1}$) as in the previous case, and we conclude as before that $a_{ij}^{q+1,1}, a_{ij}^{q+1,2}, a_{ji}^{q+1,1}, a_{ji}^{q+1,2}$ satisfy $(*)$.

Finally suppose $\{i, j\} = \{k_1, k_2\}$. We will show that $a_{k_1 k_2}^{q+1,1}$ satisfies $(*)$, whence by Lemma A.5 all four crossings of the form $a_{ij}^{q+1,l}$ for $\{i, j\} = \{k_1, k_2\}$ and $l = 1, 2$ satisfy $(*)$, and the induction step will be complete. Since Λ has at least 4 strands joining left and right, there is some other strand labeled k_3 with $k_3 \neq k_1, k_2$. There are three cases depending on the position of the y coordinate of strand k_3 relative to strands k_1 and k_2 in $[x_q, x_{q+1}]$.

If k_3 lies above both k_1 and k_2 in the y direction, then consider the two triangles shown in Figure 43. For both of these triangles, one corner is at $a_{k_1 k_2}^{q+1,1}$ and the other two corners satisfy $(*)$. Since these triangles split in two a standard bigon with corners at $a_{k_3 k_2}^{q,2}$ and $a_{k_2 k_3}^{q+1,1}$, as in the $q = 1$ case one of the triangles must be standard. It follows from Lemma A.4 that $(*)$ holds for $a_{k_1 k_2}^{q+1,1}$, as desired. If k_3 lies between k_1 and k_2 , or k_3 lies below both k_1 and k_2 , entirely similar arguments using the triangles shown in Figure 44 again show that $a_{k_1 k_2}^{q+1,1}$ satisfies $(*)$, and we are done. \square




FIGURE 43. Showing that $a_{k_1 k_2}^{q+1,1}$ satisfies $(*)$.

We can now finally prove Proposition A.3.




FIGURE 44. Two more cases to show that $a_{k_1 k_2}^{q+1,1}$ satisfies (*).

Proof of Proposition A.3. Suppose $\Psi : \mathcal{A}_{\Lambda'} \rightarrow \mathcal{A}_{\Lambda'}$ is a chain map and an isomorphism. By Lemma A.6, we have u_1, \dots, u_{2s} so that (*) holds for all Reeb chords of Λ' . Now the strands $1, \dots, 2s$ are joined in pairs at the left end of Λ' , and joined in pairs again at the right end. On the left end, for $k = 1, \dots, s$, strands $2k-1$ and $2k$ are connected, and this yields an embedded disk with a single corner at $a_{2k-1, 2k}^{1,1}$, which must be a $+$ corner by Stokes. This contributes a constant (1) term to $\delta(a_{2k-1, 2k}^{1,1})$. Since Ψ is a chain map and $\Psi(a_{2k-1, 2k}^{1,1}) = u_{2k} u_{2k-1}^{-1} a_{2k-1, 2k}^{1,1}$ by (*), it follows that $u_{2k-1} = u_{2k}$.

More generally, the same argument shows that if strands i and j are joined at either end of Λ' , then $u_i = u_j$. It follows that $u_i = u_j$ whenever i and j are part of the same connected component of Λ' . Thus we may remove duplicates and rename u_1, \dots, u_{2s} as u_1, \dots, u_m , where m is the number of components of Λ' . Then (*) becomes precisely the condition for Ψ to be a link automorphism of Λ' , and we are done. \square

A.3. Proof of Proposition A.1. With the auxiliary result Proposition A.3 in hand, we next prove Proposition A.1. Suppose that Λ_+ and Λ_- are related by a saddle cobordism at a contractible crossing flanked by mini-dips, as in the statement of Proposition A.3 or the right hand side of Figure 37. We first show that the desired map Φ is indeed a chain map, and then proceed to the main proof.

Lemma A.7. *The map $\Phi : \mathcal{A}_{\Lambda_+} \rightarrow \mathcal{A}_{\Lambda_-}$ defined in Proposition A.1 is a chain map: $\Phi \circ \partial_+ = \partial_- \circ \Phi$.*

Proof. We show that $\Phi \circ \partial_+$ and $\partial_- \circ \Phi$ agree on all Reeb chords of Λ_+ . Note that $\partial_+(a) = 0$, so $\Phi(\partial_+(a)) = 0 = \partial_-(s) = \partial_-(\Phi(a))$. Also if we denote the mini-dip crossing next to a_1 by a_3 , then $\partial_+(a_1) = \partial_-(a_1) = -a_3$, so $\Phi(\partial_+(a_1)) = -\Phi(a_3) = -a_3 = \partial_-(a_1 - s^{-1}) = \partial_-(\Phi(a_1))$; similarly $\Phi(\partial_+(a_2)) = \partial_-(\Phi(a_2))$.




FIGURE 45. Labeling the strands of Λ_+ (left) and Λ_- (right) in the cobordism region, and disks that pass through the cobordism region and contribute to $\partial_+(a_i)$ and $\partial_-(a_i)$.

Now suppose that a_i is a Reeb chord of Λ_+ besides a, a_1, a_2 : we need to show that $\Phi(\partial_+(a_i)) = \partial_-(a_i)$. The disks that make up the differentials $\partial_+(a_i)$ and $\partial_-(a_i)$ are exactly the same except where they pass through the cobordism region encompassing a_1, a, a_2 . Where these disks pass through the cobordism region, there is also a precise correspondence between the disks for $\partial_+(a_i)$ and $\partial_-(a_i)$. The oriented boundary of such a disk enters the region on one of the strands on the left and exits on one of the strands on the right, or it enters on the right and exits on the left. If we label the strands as shown in Figure 45, then for instance any disk contributing to $\partial_-(a_i)$ that enters on the left on strand 1 and exits on the right on strand 2 must pass s and turn a corner at a_1 ; there are two corresponding disks contributing to $\partial_+(a_i)$ with the same enter and exit data, one of which turns no corners in the cobordism region and one of which turns corners at a and a_1 . See Figure 45; the result replaces a monomial sa_1 in $\partial_-(a_i)$ by $1 + aa_1$ in $\partial_+(a_i)$. In all, there are 8 ways to pass through the cobordism region, with resulting contributions to $\partial_{\pm}(a_i)$ as follows:

	$\partial_+(a_i)$	$\partial_-(a_i)$		$\partial_+(a_i)$	$\partial_-(a_i)$
$1 \rightarrow 1$	a	s	$1 \leftarrow 1$	$-a_1 - a_2 - a_1aa_2$	$s^{-1} - a_1sa_2$
$1 \rightarrow 2$	$1 + aa_1$	sa_1	$1 \leftarrow 2$	$1 + aa_2$	sa_2
$2 \rightarrow 1$	$1 + a_2a$	a_2s	$2 \leftarrow 1$	$1 + a_1a$	a_1s
$2 \rightarrow 2$	$a_1 + a_2 + a_2aa_1$	$-s^{-1} + a_2sa_1$	$2 \leftarrow 2$	$-a$	$-s$

Now an inspection of this table shows that each entry in the $\partial_-(a_i)$ column is obtained from the corresponding entry in the $\partial_+(a_i)$ column by replacing a, a_1, a_2 by $s, a_1 - s^{-1}, a_2 - s^{-1}$ respectively. It immediately follows that $\Phi(\partial_+(a_i)) = \partial_-(a_i)$. \square

We now have a chain map $\Phi : \mathcal{A}_{\Lambda_+} \rightarrow \mathcal{A}_{\Lambda_-}$. In order to prove Proposition A.1, we want to show that this is equal to the geometric cobordism map Φ_{L_a} up to a link automorphism. To do this, we will first localize the differentials of Λ_{\pm} by introducing slashes in the spirit of Section A.2. In what follows, we continue to refer to the small region of Λ_{\pm} containing a_1 and a_2 (and a for Λ_+) as the ‘‘cobordism region’’, outside of which Λ_+ and Λ_- coincide. We now change Λ_- by a sequence of Reidemeister II moves that avoid the cobordism region, first pulling all vertical tangencies of $\Pi_{xy}(\Lambda_-)$ left or right so that they line up vertically, then adding slashes to separate any crossings in $\Pi_{xy}(\Lambda_-)$ or $\Pi_{xz}(\Lambda_-)$ outside the cobordism region. From this we obtain a link Λ'_- , Legendrian isotopic to Λ_- , for which there are $x_0 < x_1 < \dots < x_p$ such that:

- all vertical tangencies lie on $x = x_0$ or $x = x_p$, and the number of vertical tangencies on each of these lines is at least 2;
- there is a collection of slashes in a neighborhood of $x = x_i$ for $i = 1, \dots, p - 1$;
- there is one $i \in \{1, \dots, p - 2\}$ such that $[x_i, x_{i+1}]$ contains the cobordism region, and in that interval $[x_i, x_{i+1}]$ the only crossings in either $\Pi_{xy}(\Lambda_-)$ or $\Pi_{xz}(\Lambda_-)$ are between the two strands involved in the cobordism region;
- for every other $i = 1, \dots, p - 2$, in the interval $[x_i, x_{i+1}]$ there is exactly one crossing in either $\Pi_{xy}(\Lambda_-)$ or $\Pi_{xz}(\Lambda_-)$, and no crossing in the other;
- $[x_0, x_1]$ and $[x_{p-1}, x_p]$ contain no crossings in $\Pi_{xy}(\Lambda_-)$ or $\Pi_{xz}(\Lambda_-)$.

In short, we follow the prescription from Section A.2, except that we do not separate the crossings in the cobordism region from each other.

If we follow the same sequence of Reidemeister II moves going from Λ_- to Λ'_- , but start with Λ_+ , then we obtain a Legendrian link Λ'_+ that differs from Λ'_- only in the cobordism region.

We summarize the picture as follows:

$$\begin{array}{ccc}
 \Lambda_+ & \longrightarrow & \Lambda'_+ \\
 \downarrow & & \downarrow \\
 \Lambda_- & \longrightarrow & \Lambda'_-,
 \end{array}$$

where the horizontal arrows are Legendrian isotopies given by (the same) Reidemeister II moves, and the vertical arrows are (identical) elementary saddle cobordisms. Note that the saddle cobordism between Λ_+ and Λ_- is Hamiltonian isotopic to the concatenation of the three cobordisms specified by the other three sides of the square: from top to bottom, the isotopy from Λ_+ to Λ'_+ , followed by the saddle cobordism between Λ'_+ and Λ'_- , followed by the isotopy from Λ'_- to Λ_- . By [EHK16, Kar20], the cobordism map $\Phi_{L_a} : \mathcal{A}_{\Lambda_+} \rightarrow \mathcal{A}_{\Lambda_-}$ is chain homotopy equivalent to the composition of the cobordism maps given by the three cobordisms. We will show that this composition is the map Φ from the statement of Proposition A.1.

We first consider the cobordism map $\Phi' : \mathcal{A}_{\Lambda'_+} \rightarrow \mathcal{A}_{\Lambda'_-}$. By Lemma A.7, we know of another chain map $\Phi_1 : \mathcal{A}_{\Lambda'_+} \rightarrow \mathcal{A}_{\Lambda'_-}$: this is defined by $\Phi_1(a) = s$, $\Phi_1(a_1) = a_1 - s^{-1}$, $\Phi_1(a_2) = a_2 - s^{-1}$, and $\Phi_1(a_i) = a_i$ for all other Reeb chords a_i . Since [EHK16] gives a formula for geometric cobordism maps mod 2 and this formula is especially simple in our case, we know that the geometric map Φ' agrees with Φ_1 up to signs. By replacing s by $-s$ if necessary, we can assume that $\Phi'(a) = s$.

Lemma A.8. *There is a link automorphism $\Omega : \mathcal{A}_{\Lambda'_-} \rightarrow \mathcal{A}_{\Lambda'_-}$ such that $\Phi' = \Omega \circ \Phi_1$.*

Proof. Write ∂'_+ and ∂'_- for the differentials on $\mathcal{A}_{\Lambda'_+}$ and $\mathcal{A}_{\Lambda'_-}$ respectively.

Since the terms in Φ' agree with the terms in Φ_1 up to sign, there are signs $\sigma_i \in \{\pm 1\}$ such that $\Phi'(a_1) = \sigma_1 a_1 \pm s^{-1}$, $\Phi'(a_2) = \sigma_2 a_2 \pm s^{-1}$, and $\Phi'(a_i) = \sigma_i a_i$ for all other i . In fact, because Φ' is a chain map, we must more specifically have $\Phi'(a_1) = \sigma_1 a_1 - s^{-1}$ and $\Phi'(a_2) = \sigma_2 a_2 - s^{-1}$. To see this for a_1 (with a similar argument for a_2), we use the fact that Λ'_\pm have more than 2 strands joining left and right in the x direction, as stipulated in their construction. In particular, there is a strand of Λ'_\pm that lies either above or below the cobordism region in the xy projection. Assume this strand lies above (the argument for below is very similar). The slashes from this strand on either side of the cobordism region intersect the strands from the cobordism region in a number of crossings, two of which are labeled a_3 and a_4 in Figure 46. In $\Pi_{xy}(\Lambda'_+)$, there is a standard bigon with corners at a_3 and a_4 , contributing either a_4 to $\partial'_+(a_3)$ or a_3 to $\partial'_+(a_4)$. For definiteness assume the former (the argument is same for the latter). An inspection of Figure 46 shows that $\partial'_+(a_3)$ contains the terms $\pm(1 + a_1 a) a_4$ while $\partial'_-(a_3)$ contains $\pm a_1 s a_4$, and furthermore that these are the only terms in $\partial'_\pm(a_3)$ that involve a_4 . Since $\partial_- \Phi'(a_3) = \Phi' \partial_+(a_3)$, we must have $\pm a_1 s a_4 = \Phi'((1 + a_1 a) a_4) = \pm(1 + (\sigma_1 a_1 \pm s^{-1}) s) a_4$, which implies that the \pm sign is $-$ as claimed.




FIGURE 46. Splashes on either side of the cobordism region in Λ'_+ (left) and Λ'_- (right), with relevant crossings labeled.

Now let Ω to be the algebra automorphism of $\mathcal{A}_{\Lambda'_-}$ defined by $\Omega(a_i) = \sigma_i a_i$ for all i ; then by our expression for Φ' , we have $\Phi' = \Omega \circ \Phi_1$. It follows from the fact that Φ' and Φ_1 are both chain maps that Ω is also a chain map. Indeed, for any i we have

$$\partial'_- \Omega a_i = \partial'_- \Phi' a_i = \Phi' \partial'_+ a_i = \Omega \Phi_1 \partial'_+ a_i = \Omega \partial'_- \Phi_1 a_i = \Omega \partial'_- a_i,$$

where when $i = 1, 2$ the first and last equality follow from the fact that $\partial'_-(s^{-1}) = 0$.

It remains to show that Ω is a link automorphism of Λ'_- . To do this, we use the fact that Ω is a diagonal automorphism of $\mathcal{A}_{\Lambda'_-}$ and a chain map, and appeal to a variant of Proposition A.3. We cannot use Proposition A.3 directly because Λ'_- does not have a splash between a_1 and a_2 . However, we can still follow the inductive proof of Proposition A.3 in this setting. The only thing we need to check is the inductive step where we are given that a_1 satisfies the condition (*) from the proof and need to conclude that a_2 also satisfies this condition. To do this, let a_3 and a_5 be the crossings depicted in Figure 46, and note that there is a standard bigon in $\Pi_{xy}(\Lambda'_-)$ with corners at a_3 and a_5 . If the positive corner of this bigon is at a_3 , then $\partial'_-(a_3)$ contains the terms $(a_1 s a_2 - s^{-1}) a_5$, while if the positive corner is at a_5 , then $\partial'_-(a_5)$ contains the terms $a_3 (a_1 s a_2 - s^{-1})$. In either case, since Ω is a chain map, $\Omega(a_1) s \Omega(a_2) - s^{-1}$ must be equal to $\pm (a_1 s a_2 - s^{-1})$. Since a_1 satisfies (*), $\Omega(a_1) = u_{r(a_1)} u_{c(a_1)}^{-1} a_1$; but this implies that $\Omega(a_2) = (u_{r(a_1)} u_{c(a_1)}^{-1})^{-1} a_2 = u_{r(a_2)} u_{c(a_2)}^{-1} a_2$ and so a_2 satisfies (*), as desired. This completes the proof of Lemma A.8. \square

We next examine the maps given by the Legendrian isotopies between Λ_+ and Λ'_+ , and between Λ_- and Λ'_- . Suppose that Λ'_- is obtained from Λ_- by N Reidemeister II moves. Then we can follow [Che02, ENS02] to construct a DGA isomorphism Ψ_- between $\mathcal{A}_{\Lambda'_-}$ and the DGA $S^N(\mathcal{A}_{\Lambda_-})$ given by stabilizing \mathcal{A}_{Λ_-} N times (adding $2N$ generators in the process). This isomorphism comes from N applications of the isomorphism coming from a single Reidemeister II move, as already described in Section 4.1. By that construction, if we start with Λ_- and add the Reidemeister II moves one by one, we see that the nontrivial parts of Ψ_- come from disks with two positive punctures, one of which is at a crossing in the Reidemeister II move. By inspection, there is no point at which there is such a disk where the other positive puncture is at either a_1 or a_2 , and it follows that $\Psi_-(a_1) = a_1$ and $\Psi_-(a_2) = a_2$.

Similarly, since Λ'_+ is obtained from Λ_+ by the same Reidemeister II moves, we have a DGA isomorphism Ψ_+ between $\mathcal{A}_{\Lambda'_+}$ and $S^N(\mathcal{A}_{\Lambda_+})$, and $\Psi_+(a_1) = a_1$, $\Psi_+(a) = a$, $\Psi_+(a_2) = a_2$. Indeed, we can say more about the relation between Ψ_+ and Ψ_- . The key point is that there is a precise correspondence between the twice-positive-punctured disks that determine Ψ_+ and the twice-positive-punctured disks that determine Ψ_- : algebraically, one obtains the latter from the former by replacing a, a_1, a_2 by $s, a_1 - s^{-1}, a_2 - s^{-1}$ just as in the proof of Lemma A.7. Consequently, for any Reeb chord a_i of Λ'_+ (and thus of Λ'_-) besides a, a_1, a_2 , $\Psi_-(a_i)$ is obtained from $\Psi_+(a_i)$ by this algebraic replacement.

Put another way, let Φ_1 be as above, and similarly define $\Phi_2 : S^N(\mathcal{A}_{\Lambda_+}) \rightarrow S^N(\mathcal{A}_{\Lambda_-})$ by $\Phi_2(a) = s$, $\Phi_2(a_1) = a_1 - s^{-1}$, $\Phi_2(a_2) = a_2 - s^{-1}$, and Φ_2 is the identity on all other generators of $S^N(\mathcal{A}_{\Lambda_+})$. Note that by Lemma A.7, Φ_1 and Φ_2 are both chain maps. By the above discussion, we conclude that the following diagram commutes:

$$\begin{array}{ccc} S^N(\mathcal{A}_{\Lambda_+}) & \xleftarrow[\cong]{\Psi_+} & \mathcal{A}_{\Lambda'_+} \\ \downarrow \Phi_2 & & \downarrow \Phi_1 \\ S^N(\mathcal{A}_{\Lambda_-}) & \xleftarrow[\cong]{\Psi_-} & \mathcal{A}_{\Lambda'_-} \end{array}$$

The cobordism map $\mathcal{A}_{\Lambda_+} \rightarrow \mathcal{A}_{\Lambda'_+}$ is simply the composition of the inclusion map $i : \mathcal{A}_{\Lambda_+} \rightarrow S^N(\mathcal{A}_{\Lambda_+})$ and the inverse of Ψ_+ , and the cobordism map $\mathcal{A}_{\Lambda'_-} \rightarrow \mathcal{A}_{\Lambda_-}$ is the composition of Ψ_- and the projection map $p : S^N(\mathcal{A}_{\Lambda_-}) \rightarrow \mathcal{A}_{\Lambda_-}$.

We can now finally turn to the geometric cobordism map $\Phi_{L_a} : \mathcal{A}_{\Lambda_+} \rightarrow \mathcal{A}_{\Lambda_-}$. To complete the proof of Proposition A.1, we want to show that $\Phi_{L_a} = \Omega \circ \Phi$ for some link automorphism Ω of Λ_- .

At this point we have broken down Φ_{L_a} into a composition of three cobordism maps: $\Psi_+^{-1} \circ i : \mathcal{A}_{\Lambda_+} \rightarrow \mathcal{A}_{\Lambda'_+}$, $\Phi' : \mathcal{A}_{\Lambda'_+} \rightarrow \mathcal{A}_{\Lambda'_-}$, and $p \circ \Psi_- : \mathcal{A}_{\Lambda'_-} \rightarrow \mathcal{A}_{\Lambda_-}$. That is, Φ_{L_a} is chain homotopy equivalent to the composition $p \circ \Psi_- \circ \Phi' \circ \Psi_+^{-1} \circ i$ of the five maps going around the sides of the following rectangle:

$$\begin{array}{ccccc}
 \mathcal{A}_{\Lambda_+} & \xrightarrow{i} & S^N(\mathcal{A}_{\Lambda_+}) & \xleftarrow[\cong]{\Psi_+} & \mathcal{A}_{\Lambda'_+} \\
 \downarrow \Phi_{L_a} & & \downarrow \Phi_2 & & \downarrow \Phi' \\
 \mathcal{A}_{\Lambda_-} & \xleftarrow[p]{} & S^N(\mathcal{A}_{\Lambda_-}) & \xleftarrow[\cong]{\Psi_-} & \mathcal{A}_{\Lambda'_-}
 \end{array}$$

From Lemma A.8, there is a link automorphism Ω of Λ'_- such that $\Phi' = \Omega \circ \Phi_1$. Since Λ'_- and Λ_- are Legendrian isotopic, Ω induces a link automorphism of Λ_- , which we also call Ω , so that Ω commutes with the chain map $p \circ \Psi_- : \mathcal{A}_{\Lambda'_-} \rightarrow \mathcal{A}_{\Lambda_-}$ induced by the isotopy. Thus

$$\Phi_{L_a} \simeq p \circ \Psi_- \circ \Phi' \circ \Psi_+^{-1} \circ i = p \circ \Psi_- \circ \Omega \circ \Phi_1 \circ \Psi_+^{-1} \circ i = \Omega \circ p \circ \Psi_- \circ \Phi_1 \circ \Psi_+^{-1} \circ i = \Omega \circ p \circ \Phi_2 \circ i.$$

But $p \circ \Phi_2 \circ i$ is exactly equal to Φ as defined in the statement of Proposition A.1, and we are done with the proof.

A.4. Proof of Proposition 4.8. The remainder of this section is devoted to the proof of Proposition 4.8. At this point, by Proposition A.1, we know the saddle cobordism map for a saddle flanked by mini-dips; to prove Proposition 4.8, we just need to compose this map with maps corresponding to the Reidemeister II moves of adding and removing mini-dips. This is similar to the proof of Proposition A.1 in the previous subsection, except that it will now be important to calculate these Reidemeister II maps in more detail.

Suppose that, as in the statement of Proposition 4.8, we have a saddle cobordism between Λ_+ and Λ_- , where the cobordism is given by resolving a proper contractible Reeb chord a of Λ_+ . Let Λ'_+ be the result of adding a mini-dip to Λ_+ just after a following the orientation of Λ_+ , and let Λ''_+ be result of further adding a mini-dip to Λ'_+ on the other side of a . Similarly define Λ'_- and Λ''_- . Then Λ'_\pm are obtained from Λ_\pm by a single Reidemeister II move, Λ''_\pm are obtained from Λ'_\pm by another Reidemeister II move, and Λ''_+ and Λ''_- are related by a saddle move of the precise form that we considered in Proposition A.1. See Figure 47.




FIGURE 47. Adding mini-dips to Λ_\pm .

The properness condition for a translates into the following result.

Lemma A.9. *Given that the Reeb chord a is proper:*

- if a_q is any Reeb chord of Λ'_+ , and ∂' denotes the differential on Λ'_+ , then any term in $\partial'(a_q)$ that contains a_3 must contain a_3 exactly once and cannot contain a_1 ;
- if a_q is any Reeb chord of Λ''_+ , and ∂'' denotes the differential on Λ''_+ , then any term in $\partial''(a_q)$ that contains a_4 must contain a_4 exactly once and cannot contain any of a_1, a_2, a_3 .




FIGURE 48. Turning an immersed disk Δ' for Λ'_+ into an immersed disk Δ for Λ_+ .

Proof. We will establish the statement for Λ'_+ ; the proof of the statement for Λ''_+ is similar. If a_q is any of a, a_1, a_3 , then the statement is trivially true: by action considerations, the only term in $\partial'(a_q)$ that could contain a_3 is just the term a_3 itself in $\partial'(a_1)$. Now assume a_q is not a, a_1, a_3 . Consider any word in $\partial'(a_q)$, corresponding to an immersed disk Δ' in Λ'_+ with sole $+$ corner at a_q and a $-$ corner at a_3 . Then Δ' in turn produces an immersed disk Δ in Λ_+ , now possibly with concave corners at a : see Figure 48. If Δ' contained multiple corners at a_3 , or corners at both a_1 and a_3 , then the boundary of Δ would pass through $\Pi_{xy}(a)$ more than once, violating the properness condition from Definition 4.3. \square

We will now piece together the five maps $\mathcal{A}_{\Lambda_+} \rightarrow \mathcal{A}_{\Lambda'_+} \rightarrow \mathcal{A}_{\Lambda''_+} \rightarrow \mathcal{A}_{\Lambda''_-} \rightarrow \mathcal{A}_{\Lambda'_-} \rightarrow \mathcal{A}_{\Lambda_-}$ to get the desired cobordism map. The central map $\mathcal{A}_{\Lambda''_+} \rightarrow \mathcal{A}_{\Lambda''_-}$ has already been computed, while the remaining maps come from Reidemeister II isotopies.

We will focus for now on the map $\mathcal{A}_{\Lambda_+} \rightarrow \mathcal{A}_{\Lambda'_+}$, which we call Ψ_+^{\rightarrow} . This is the chain map induced by adding a Legendrian Reidemeister II move, as derived in [Che02, ENS02] and summarized in Section 4.1 above, and we describe it explicitly now. Label the Reeb chords of Λ_+ besides a as a_5, \dots, a_r , so that we can write $\mathcal{A}_{\Lambda_+} = \mathcal{A}(a, a_5, \dots, a_r)$ and $\mathcal{A}_{\Lambda'_+} = \mathcal{A}(a, a_1, a_3, a_5, \dots, a_r)$. We stabilize \mathcal{A}_{Λ_+} by adding two new generators e_1, e_2 with $|e_1| = 0$, $|e_2| = -1$, $\partial(e_1) = e_2$, $\partial(e_2) = 0$, to produce a new DGA $S(\mathcal{A}_{\Lambda_+}) = \mathcal{A}(a, a_5, \dots, a_r, e_1, e_2)$. As described in Section 4.1 and specifically defined in (4.1), there is a chain isomorphism $\Psi : \mathcal{A}_{\Lambda'_+} \rightarrow S(\mathcal{A}_{\Lambda_+})$, which in our case is defined by $\Psi(a_1) = e_1$, $\Psi(a_3) = -e_2$, $\Psi(a) = a$, and for $\ell \geq 5$,

$$\Psi(a_\ell) = a_\ell - H\Psi\partial'a_\ell$$

where ∂' is the differential on Λ'_+ . Then Ψ_+^{\rightarrow} is defined to be equal to $\Psi^{-1} \circ i$.

We now claim that Ψ_+^{\rightarrow} satisfies the following formula, which can be compared to the definition of Φ^{\rightarrow} from Section 4.2.

Lemma A.10. *For all $\ell \geq 5$, we have*

$$(A.1) \quad \Psi_+^{\rightarrow}(a_\ell) = a_\ell - \sum_{\Delta \in \Delta_a^{\rightarrow}(a_\ell)} (-1)^{|w_1(\Delta)|} \text{sgn}(\Delta) \Psi_+^{\rightarrow}(w_1(\Delta)) a_1 w_2(\Delta).$$

Proof. Assume without loss of generality that a, a_5, \dots, a_r are ordered by height (note that a is contractible and thus has the shortest height). We first claim that for all $q \geq 5$, $\Psi(a_q) - a_q$ only includes terms that involve at least one e_1 and no e_2 : we abbreviate this condition by $\Psi(a_q) - a_q = O(e_1)$. We prove this by induction on q , where the base case is actually $a_q = a$ (note $\Psi(a) - a = 0$). For the induction step, note that $\Psi(a_q) - a_q = -H\Psi\partial'a_q$, and the right hand side only contains terms involving at least one e_1 ; we need to show that $H\Psi\partial'a_q$ does not involve e_2 .

Consider any word w in $\partial'a_q$. If a_3 does not appear in w , then w involves only $a, a_1, a_5, \dots, a_{q-1}$, and so by induction $\Psi(w) - w = O(e_1)$ and $H\Psi(w) = H(w) = 0$. On the other hand, if w does involve a_3 , then by Lemma A.9, $w = w_1 a_3 w_2$ where w_1, w_2 involve only a, a_5, \dots, a_{q-1} ; then by induction again, $H\Psi(w) = -H((\Psi(w_1))e_2(\Psi(w_2))) = -H(w_1 e_2 \Psi(w_2)) = \pm w_1 e_1 \Psi(w_2)$ does not involve e_2 . This completes the proof that $\Psi(a_q) - a_q = O(e_1)$ for all $q \geq 5$.

We now prove the lemma, again by induction on ℓ . The base case is actually $\Psi_+^\rightarrow(a) = a$, which is (A.1) with $a = a_\ell$. For the induction step, we compute that:

$$\Psi_+^\rightarrow(a_\ell) = \Psi^{-1}(a_\ell) = a_\ell + \Psi^{-1}H\Psi\partial'a_\ell.$$

Now suppose that w is a word in $\partial'a_\ell$, and again apply Lemma A.9. If w does not contain a_3 , then $H\Psi(w) = 0$. If w does contain a_3 , then we write $w = w_1 a_3 w_2$ and compute:

$$H\Psi(w) = H\Psi(w_1 a_3 w_2) = -H(\Psi(w_1)e_2\Psi(w_2)) = -H(w_1 e_2 \Psi(w_2)) = (-1)^{|w_1|} w_1 e_1 \Psi(w_2)$$

and thus

$$\Psi^{-1}H\Psi(w) = (-1)^{|w_1|} \Psi^{-1}(w_1 e_1 \Psi(w_2)) = (-1)^{|w_1|} \Psi^{-1}(w_1) a_1 w_2 = (-1)^{|w_1|} \Psi_+^\rightarrow(w_1) a_1 w_2,$$

where we have used the fact that w_1 does not involve a_1 or a_3 and thus $\Psi^{-1}(w_1) = \Psi^{-1}i(w_1) = \Psi_+^\rightarrow(w_1)$. Finally note that the disk for w in Λ'_+ precisely corresponds to a disk Δ in $\Delta_a^\rightarrow(a_\ell)$ in Λ_+ , and that the sign for w in $\partial'a_\ell$ is $-\text{sgn}(\Delta)$ since Δ replaces a corner at a_3 with positive orientation sign with a corner at a with negative orientation sign. Now the signed sum of $\Psi^{-1}H\Psi(w) = \Psi_+^\rightarrow(w_1) a_1 w_2$ over all disks in $\Delta_a^\rightarrow(a_\ell)$ gives (A.1), and this completes the induction. \square

In a similar way, we write Ψ_+^\leftarrow for the cobordism map from $\mathcal{A}_{\Lambda'_+} = \mathcal{A}(a, a_1, a_3, a_5, \dots, a_r)$ to $\mathcal{A}_{\Lambda''_+} = \mathcal{A}(a, a_1, a_2, a_3, a_4, a_5, \dots, a_r)$ induced by the Reidemeister II isotopy between Λ'_+ and Λ''_+ .

Lemma A.11. *For all $\ell \geq 5$, we have*

$$\Psi_+^\leftarrow(a_\ell) = a_\ell - \sum_{\Delta \in \Delta_a^\leftarrow(a_\ell)} (-1)^{|w_1(\Delta)|} \text{sgn}(\Delta) \Psi_+^\rightarrow(w_1(\Delta)) a_2 w_2(\Delta).$$

Proof. This is essentially identical to the proof of Lemma A.10. Given our choice of orientation signs, there are two sign differences here from the proof of Lemma A.10: $\Psi(a_4)$ is now e_2 rather than $-e_2$, and the sign of a word contributing to $\partial'(a_\ell)$ is now equal to $+\text{sgn}(\Delta)$ rather than $-\text{sgn}(\Delta)$ for the corresponding disk $\Delta \in \Delta_a^\leftarrow(a_\ell)$. These two sign changes cancel out. One other subtle difference is that if we follow the proof of the previous lemma, then $\Delta_a^\leftarrow(a_\ell)$ in the statement of the present lemma should be for Λ'_+ rather than for Λ_+ . However, by the properness condition for a , there is a one-to-one correspondence between disks in $\Delta_a^\leftarrow(a_\ell)$ for Λ'_+ and Λ_+ , and so the desired formula holds for either form of $\Delta_a^\leftarrow(a_\ell)$. \square

We can now finally piece together our various subsidiary results to prove Proposition 4.8. To distinguish between the saddle cobordisms in the dipped and undipped settings, let L_a be the cobordism between Λ_+ and Λ_- as in the statement of Proposition 4.8, and let \tilde{L}_a be the cobordism between Λ''_+ and Λ''_- . As shown in Figure 47, we can concatenate \tilde{L}_a and four Lagrangians coming from Legendrian isotopies to create a five-story cobordism between Λ_+

and Λ_- which is Hamiltonian isotopic to L_a : from top to bottom, the five cobordisms go between Λ_+ , Λ'_+ , Λ''_+ , Λ'_- , and Λ_- .

The chain map $\Phi_{L_a} : \mathcal{A}_{\Lambda_+} \rightarrow \mathcal{A}_{\Lambda_-}$ is then chain homotopic to the composition of the chain maps coming from the five cobordisms. We summarize this in the following diagram, which commutes up to chain homotopy:

$$\begin{array}{ccccc}
\mathcal{A}_{\Lambda_+} & \xrightarrow{\Psi_+^\rightarrow} & \mathcal{A}_{\Lambda'_+} & \xrightarrow{\Psi_+^\leftarrow} & \mathcal{A}_{\Lambda''_+} \\
\downarrow \Phi_{L_a} & & & & \downarrow \Phi_{\tilde{L}_a} \\
\mathcal{A}_{\Lambda_-} & \xleftarrow{p_2} & \mathcal{A}_{\Lambda'_-} & \xleftarrow{p_1} & \mathcal{A}_{\Lambda''_-} .
\end{array}$$

Here Ψ_+^\rightarrow and Ψ_+^\leftarrow are the maps computed in Lemmas A.10 and A.11, while p_1 and p_2 are the maps induced by the reverse Reidemeister II moves from Λ''_- to Λ'_- and from Λ'_- to Λ_- . By Remark 4.2, these last two maps (which correspond to $p \circ \Psi$ in Remark 4.2) are given simply by projection: $p_1(a_2) = p_1(a_4) = p_2(a_1) = p_2(a_3) = 0$ and p_1, p_2 are the identity on all other generators.

Now by Proposition A.1, $\Phi_{\tilde{L}_a} = \Omega \circ \Phi$ where Ω is a link automorphism of Λ''_- and Φ is the map given in the statement of the proposition. Since Λ''_- and Λ_- are isotopic, Ω induces a link automorphism of Λ_- which we also denote by Ω , and $p_2 \circ p_1 \circ \Omega = \Omega \circ p_2 \circ p_1$. At this point we have:

$$\Phi_{L_a} \simeq p_2 \circ p_1 \circ \Phi_{\tilde{L}_a} \circ \Psi_+^\leftarrow \circ \Psi_+^\rightarrow = \Omega \circ p_2 \circ p_1 \circ \Phi \circ \Psi_+^\leftarrow \circ \Psi_+^\rightarrow.$$

We will be done if we can show that the composition $p_2 \circ p_1 \circ \Phi \circ \Psi_+^\leftarrow \circ \Psi_+^\rightarrow$ is equal to the map $\Phi_{L_a}^{\text{comb}} = \Phi^\leftarrow \circ \Phi^\rightarrow \circ \Phi_0$ from Proposition 4.8. But $\Phi_{L_a}^{\text{comb}}$ is specifically designed so that this is the case. Specifically, if a_ℓ is any Reeb chord of Λ_+ besides a , then $\Phi^\rightarrow(a_\ell)$ and $\Phi^\leftarrow(a_\ell)$ are precisely the result of replacing a_1 and a_2 by $-s^{-1}$ in the expressions for $\Psi_+^\rightarrow(a_\ell)$ and $\Psi_+^\leftarrow(a_\ell)$ from Lemmas A.10 and A.11. But by the definition of Φ , this replacement is exactly the effect of composing with the map $p_2 \circ p_1 \circ \Phi$, which sends a_1, a_2 to $-s^{-1}$ and sends a_ℓ to itself for $\ell \geq 5$. It follows that

$$\Phi_{L_a}^{\text{comb}}(a_\ell) = (\Phi^\leftarrow \circ \Phi^\rightarrow)(a_\ell) = (p_2 \circ p_1 \circ \Phi)((\Psi_+^\leftarrow \circ \Psi_+^\rightarrow)(a_\ell))$$

for all ℓ . Combined with the fact that $\Phi_{L_a}^{\text{comb}}(a) = s = (p_2 \circ p_1 \circ \Phi \circ \Psi_+^\leftarrow \circ \Psi_+^\rightarrow)(a)$, this establishes that $\Phi_{L_a}^{\text{comb}} = p_2 \circ p_1 \circ \Phi \circ \Psi_+^\leftarrow \circ \Psi_+^\rightarrow$. The proof of Proposition 4.8 is complete.

REFERENCES

- [Abb14] Casim Abbas. *An introduction to compactness results in symplectic field theory*. Springer, Heidelberg, 2014.
- [Ad90] V. I. Arnol' d. *Singularities of caustics and wave fronts*, volume 62 of *Mathematics and its Applications (Soviet Series)*. Kluwer Academic Publishers Group, Dordrecht, 1990.
- [BEH⁺03] F. Bourgeois, Y. Eliashberg, H. Hofer, K. Wysocki, and E. Zehnder. Compactness results in symplectic field theory. *Geom. Topol.*, 7:799–888, 2003.
- [Ben83] Daniel Bennequin. Entrelacements et équations de Pfaff. In *Third Schnepfenried geometry conference, Vol. 1 (Schnepfenried, 1982)*, volume 107 of *Astérisque*, pages 87–161. Soc. Math. France, Paris, 1983.
- [BFZ05] Arkady Berenstein, Sergey Fomin, and Andrei Zelevinsky. Cluster algebras. III. Upper bounds and double Bruhat cells. *Duke Math. J.*, 126(1):1–52, 2005.
- [BLL18] Sebastian Baader, Lukas Lewark, and Livio Liechti. Checkerboard graph monodromies. *Enseign. Math.*, 64(1-2):65–88, 2018.
- [Cas21] Roger Casals. Lagrangian skeleta and plane curve singularities. *J. Fixed Point Theory and Applications (Viterbo 60 Volume)*, 2021.
- [CDGG17] Baptiste Chantraine, Georgios Dimitroglou Rizell, Paolo Ghiggini, and Roman Golovko. Geometric generation of the wrapped Fukaya category of Weinstein manifolds and sectors. arXiv:1712.09126, 2017.

- [CDGG20] Baptiste Chantraine, Georgios Dimitroglou Rizell, Paolo Ghiggini, and Roman Golovko. Floer theory for Lagrangian cobordisms. *J. Differential Geom.*, 114(3):393–465, 2020.
- [CG21] Roger Casals and Honghao Gao. Infinitely many Lagrangian fillings. *Annals of Math.*, 2021.
- [CGGS20] Roger Casals, Eugene Gorsky, Mikhail Gorsky, and José Simental. Algebraic weaves and braid varieties. arXiv:2012.06931, 2020.
- [Cha10] Baptiste Chantraine. Lagrangian concordance of Legendrian knots. *Algebr. Geom. Topol.*, 10(1):63–85, 2010.
- [Che02] Yuri Chekanov. Differential algebra of Legendrian links. *Invent. Math.*, 150(3):441–483, 2002.
- [CSLL⁺21] Orsola Capovilla-Searle, Noémie Legout, Maÿlys Limouzeineau, Emmy Murphy, Yu Pan, and Lisa Traynor. Obstructions to exact Lagrangian cobordisms. In preparation, 2021.
- [CZ21] Roger Casals and Eric Zaslow. Legendrian weaves. To appear *Geometry & Topology*, 2021.
- [EES05] Tobias Ekholm, John Etnyre, and Michael Sullivan. Orientations in Legendrian contact homology and exact Lagrangian immersions. *Internat. J. Math.*, 16(5):453–532, 2005.
- [EHK16] Tobias Ekholm, Ko Honda, and Tamás Kálmán. Legendrian knots and exact Lagrangian cobordisms. *J. Eur. Math. Soc. (JEMS)*, 18(11):2627–2689, 2016.
- [EK08] Tobias Ekholm and Tamás Kálmán. Isotopies of Legendrian 1-knots and Legendrian 2-tori. *J. Symplectic Geom.*, 6(4):407–460, 2008.
- [Ekh19] Tobias Ekholm. Holomorphic curves for legendrian surgery. arXiv:1906.07228, 2019.
- [EL19] Tobias Ekholm and Yanki Lekili. Duality between lagrangian and legendrian invariants. arXiv:1701.01284, 2019.
- [EN18] John B. Etnyre and Lenhard L. Ng. Legendrian contact homology in \mathbb{R}^3 . arXiv:1811.10966, 2018.
- [ENS02] John B. Etnyre, Lenhard L. Ng, and Joshua M. Sabloff. Invariants of Legendrian knots and coherent orientations. *J. Symplectic Geom.*, 1(2):321–367, 2002.
- [EP96] Y. Eliashberg and L. Polterovich. Local Lagrangian 2-knots are trivial. *Ann. of Math. (2)*, 144(1):61–76, 1996.
- [Etn03] John B. Etnyre. Introductory lectures on contact geometry. In *Topology and geometry of manifolds (Athens, GA, 2001)*, volume 71 of *Proc. Sympos. Pure Math.*, pages 81–107. Amer. Math. Soc., Providence, RI, 2003.
- [EV18] John Etnyre and Vera Vértési. Legendrian satellites. *Int. Math. Res. Not. IMRN*, (23):7241–7304, 2018.
- [FR11] Dmitry Fuchs and Dan Rutherford. Generating families and Legendrian contact homology in the standard contact space. *J. Topol.*, 4(1):190–226, 2011.
- [Gei08] Hansjörg Geiges. *An introduction to contact topology*, volume 109 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2008.
- [GPS19] Sheel Ganatra, John Pardon, and Vivek Shende. Sectorial descent for wrapped Fukaya categories. arXiv:1809.03427, 2019.
- [GSW20a] Honghao Gao, Linhui Shen, and Daping Weng. Augmentations, fillings, and clusters. arXiv:2008.10793, 2020.
- [GSW20b] Honghao Gao, Linhui Shen, and Daping Weng. Positive braid links with infinitely many fillings. arXiv:2009.00499, 2020.
- [Kál05] Tamás Kálmán. Contact homology and one parameter families of Legendrian knots. *Geom. Topol.*, 9:2013–2078, 2005.
- [Kál06] Tamás Kálmán. Braid-positive Legendrian links. *Int. Math. Res. Not.*, pages Art ID 14874, 29, 2006.
- [Kar17] Cecilia Karlsson. To compute orientations of Morse flow trees in Legendrian contact homology. arXiv:1704.05156, 2017.
- [Kar20] Cecilia Karlsson. A note on coherent orientations for exact Lagrangian cobordisms. *Quantum Topol.*, 11(1):1–54, 2020.
- [Lev16] C. Levenson. Augmentations and rulings of Legendrian knots. *J. Symplectic Geom.*, 14(4):1089–1143, 2016.
- [Lev17] Caitlin Levenson. Augmentations and rulings of Legendrian links in $\#^k(S^1 \times S^2)$. *Pacific J. Math.*, 288(2):381–423, 2017.
- [Lin16] Francesco Lin. Exact Lagrangian caps of Legendrian knots. *J. Symplectic Geom.*, 14(1):269–295, 2016.
- [Mis03] K. Mishachev. The N -copy of a topologically trivial Legendrian knot. *J. Symplectic Geom.*, 1(4):659–682, 2003.
- [Mis17] Filip Misev. Cutting arcs for torus links and trees. *Bull. Soc. Math. France*, 145(3):575–602, 2017.
- [Ng] Lenhard Ng. *Mathematica notebook fillings.nb*. Available at <https://math.duke.edu/~ng/math/>.
- [Ng03] Lenhard L. Ng. Computable Legendrian invariants. *Topology*, 42(1):55–82, 2003.
- [Ng10] Lenhard Ng. Rational symplectic field theory for Legendrian knots. *Invent. Math.*, 182(3):451–512, 2010.

- [NR13] Lenhard Ng and Daniel Rutherford. Satellites of Legendrian knots and representations of the Chekanov-Eliashberg algebra. *Algebr. Geom. Topol.*, 13(5):3047–3097, 2013.
- [NRS⁺20] Lenhard Ng, Dan Rutherford, Vivek Shende, Steven Sivek, and Eric Zaslow. Augmentations are Sheaves. *Geom. Topol.*, 24(5):2149–2286, 2020.
- [NT04] Lenhard Ng and Lisa Traynor. Legendrian solid-torus links. *J. Symplectic Geom.*, 2(3):411–443, 2004.
- [Oh15] Yong-Geun Oh. *Symplectic topology and Floer homology. Vol. 1*, volume 28 of *New Mathematical Monographs*. Cambridge University Press, Cambridge, 2015. Symplectic geometry and pseudo-holomorphic curves.
- [OS04] Burak Ozbagci and András I. Stipsicz. *Surgery on contact 3-manifolds and Stein surfaces*, volume 13 of *Bolyai Society Mathematical Studies*. Springer-Verlag, Berlin, 2004.
- [Pan17a] Yu Pan. The augmentation category map induced by exact Lagrangian cobordisms. *Algebr. Geom. Topol.*, 17(3):1813–1870, 2017.
- [Pan17b] Yu Pan. Exact Lagrangian fillings of Legendrian $(2, n)$ torus links. *Pacific J. Math.*, 289(2):417–441, 2017.
- [Pol01] Leonid Polterovich. *The geometry of the group of symplectic diffeomorphisms*. Lectures in Mathematics ETH Zürich. Birkhäuser Verlag, Basel, 2001.
- [Rud92] Lee Rudolph. Quasipositive annuli. (Constructions of quasipositive knots and links. IV). *J. Knot Theory Ramifications*, 1(4):451–466, 1992.
- [Sab05] Joshua M. Sabloff. Augmentations and rulings of Legendrian knots. *Int. Math. Res. Not.*, (19):1157–1180, 2005.
- [Sau04] Denis Sauvaget. Curiosités lagrangiennes en dimension 4. *Ann. Inst. Fourier (Grenoble)*, 54(6):1997–2020 (2005), 2004.
- [Siv11] Steven Sivek. A bordered Chekanov-Eliashberg algebra. *J. Topol.*, 4(1):73–104, 2011.

UNIVERSITY OF CALIFORNIA DAVIS, DEPT. OF MATHEMATICS, SHIELDS AVENUE, DAVIS, CA 95616, USA
E-mail address: `casals@math.ucdavis.edu`

DUKE UNIVERSITY, DEPARTMENT OF MATHEMATICS, DURHAM, NC 27708, USA
E-mail address: `ng@math.duke.edu`