THE ACYCLICITY OF THE COMPLEX OF HOMOLOGOUS CURVES

DANIEL MINAHAN

Abstract. We show that the complex of homologous curves of a surface of genus g is $(g-3)$–acyclic.

1. Introduction

Let S_g be a closed, oriented surface of genus g. Let $\vec{x} \in H_1(S_g;\mathbb{Z})$ be a nonzero primitive homology class. The *complex of homologous curves* $\mathcal{C}_{\vec{x}}(S_g)$, defined by Putman [Put08], is the complex where a k–cell is an oriented multicurve $M = x_0 \sqcup \ldots \sqcup x_k$ such that $[x_i] = \vec{x}$ for every $0 \leq i \leq k$, where $[x_i]$ denotes the homology class of x_i. Putman used the work of Johnson [Joh83] to show that when $g \geq 3$, the complex $C_{\vec{x}}(S_g)$ is connected [Put12]. The main goal of the paper is to prove the following theorem.

Theorem A. *Let $g \geq 2$. The integral homology $\widetilde{H}_k(\mathcal{C}_{\vec{x}}(S_g);\mathbb{Z})$ vanishes for $k \leq g-3$.*

Application of $C_{\vec{x}}(S_g)$ to the Torelli group. The *Torelli group*, denoted $\mathcal{I}_g$, is the subgroup of the mapping class group consisting of mapping classes φ that act trivially on $H_1(S_g;\mathbb{Z})$. The Torelli group $\mathcal{I}_g$ acts naturally on $C_{\vec{x}}(S_g)$ for any choice of primitive nonzero $\vec{x} \in H_1(S_g;\mathbb{Z})$. The complex $C_{\vec{x}}(S_g)$ has been used by Hatcher and Margalit to give a new proof that $\mathcal{I}_g$ is generated by bounding pair maps [HM12]. Gaster, Greene and Vlamis also connected colorings of $C_{\vec{x}}(S_g)$ with the Chillingworth homomorphism [GGV18].

1.1. **The strategy of the proof of Theorem A.** Let $g \geq 2$ and let $\vec{x} \in H_1(S_g;\mathbb{Z})$ be a nonzero primitive homology class. Bestvina, Bux and Margalit defined a complex called the complex of minimizing cycles, denoted $\mathcal{B}_{\vec{x}}(S_g)$ [BBM10]. The complex of minimizing cycles has two important properties for our purposes:

- $C_{\vec{x}}(S_g)$ is a subcomplex of $\mathcal{B}_{\vec{x}}(S_g)$, and
- $\mathcal{B}_{\vec{x}}(S_g)$ is contractible [BBM10, Theorem E].

Date: May 7, 2023.

2000 *Mathematics Subject Classification.* Primary 54C40, 14E20; Secondary 46E25, 20C20.

Key words and phrases. Curve complex.

Hatcher and Margalit [HM12] use PL–Morse theory to prove that $H_1(\mathcal{B}_{\vec{x}}(S_g), C_{\vec{x}}(S_g); \mathbb{Z}) = 0$ when $g \geq 3$, which, along with the contractibility of $\mathcal{B}_{\vec{x}}(S_g)$, implies that $C_{\vec{x}}(S_g)$ is connected when $g \geq 3$. We will use the same PL–Morse function to prove $H_k(\mathcal{B}_{\vec{x}}(S_g)/C_{\vec{x}}(S_g); \mathbb{Z}) = 0$ for $k \leq g - 2$. This and the long exact sequence in homology immediately prove Theorem A.

As part of the proof that $H_k(\mathcal{B}_{\vec{x}}(S_g)/C_{\vec{x}}(S_g); \mathbb{Z}) = 0$ for $k \leq g-2$, we will also prove a result about the *complex of splitting curves.* Let S_g^b denote a compact oriented surface of genus g with b boundary components. We say that a curve $\delta \subseteq S_g^b$ is separating if $S_g^b \setminus \delta$ is disconnected. Let $g \geq 2$ and $S = S_g^2$. The *complex of separating curves* $C_{\text{sep}}(S)$ is the full subcomplex of $C(S)$ generated by curves δ such that $S \setminus \delta$ is disconnected. The *complex of splitting curves* $C_{\text{split}}(S_g^2)$ is the full subcomplex of $C_{\text{sep}}(S)$ generated by curves δ such that the connected components of ∂S are in different connected components of $S \setminus \delta$. We will prove in Proposition 5.1 that $C_{\text{split}}(S)$ is $(g-3)$–acyclic.

1.2. **Outline of the paper.** The paper is organized into the following chunks.

- General connectivity and acyclicity results (Sections 2, 3, and 4)
- Proof of Proposition 5.1 (Section 5)
- Proof of Theorem A (Sections 6 and 7)

We now give an overview of each section.

Section 2. We discuss some general facts about connectivity and acyclicity of simplicial complexes. We then prove Lemma 2.3, which says that the relative homology of a pair of simplicial complexes $A \subseteq B$ can be computed using PL–Morse theory. Specifically, suppose that there is a function $W : B^{(0)} \to \mathbb{Z}_{\geq 0}$ such that $W^{-1}(\{0\}) = A^{(0)}$. We will show that if the function W satisfies certain nice local acyclicity properties, then the relative homology $H_k(B, A; \mathbb{Z})$ vanishes in a range depending on W.

Section 3. We prove Proposition 3.1. This is a packaging of standard results about the Čech–to–singular spectral sequence. We will assume that we have some simplicial complex A and a simplicial cover $\mathcal{U}$ of A with $\mathcal{U}$ "indexed" by another simplicial complex B. We will show that if $\widetilde{H}_k(B; \mathbb{Z})$ vanishes in a range and the elements of the cover $\mathcal{U}$ also satisfy some acyclicity properties, then $\widetilde{H}_k(A; \mathbb{Z})$ also vanishes in a range.

Section 4. We prove Lemma 4.1. This is a result based on the work of Brendle, Broaddus and Putman [BBP23] that allows us to compute the acyclicity of certain subcomplexes of the curve complex of surfaces with boundary.

Section 5. We prove Proposition 5.1. We will use Lemma 2.3 to prove that a variant of the arc complex on surfaces with certain decoration on the boundary are acyclic in a range. We then use Proposition 3.1 to prove that a more general version of the complex of splitting curves is acyclic in a range. This will imply Proposition 5.1.

Section 6. We will revisit some of the ideas from Section 2 in a slightly different context. In particular, the results in Section 2 apply only to simplicial complexes. However, the complex of minimizing cycles $\mathcal{B}_{\vec{x}}(S_g)$ is not a simplicial complex. We will resolve this issue in Lemma 6.6 by showing that Lemma 2.3 can be applied to CW–complexes equipped with some convex structure.

Section 7. We will prove Theorem A. We will apply a variant of Lemma 2.3. The required local acyclicity properties will be verified by inductively applying Proposition 4.1. The base case of this argument uses Proposition 5.1.

Acknowledgments. We would like to thank our advisor Dan Margalit for many helpful math conversations and his continued support and encouragement. We would also like to thank Benson Farb and Andy Putman for reading and providing comments on earlier drafts of this paper. We would like to thank Igor Belegradek, Wade Bloomquist, Katherine Williams Booth, John Etnyre, Annie Holden, Thang Le, Agniva Roy, Roberta Shapiro and Cindy Tan for helpful conversations regarding this paper.

2. Connectivity and PL–Morse theory

In this section, we explain some basic algebraic topology facts. We then explain some of the basic ideas of PL-Morse theory, namely:

- PL–Morse functions and
- descending links.

For more background, see Bestvina's survey [Bes]. We begin with Section 2.1, which includes some basic definitions and results about connectivity and acyclicity. In Section 2.2, we prove Lemma 2.2, which is a slight reformulation of Bestvina's results. We then use Lemma 2.2 to prove Lemma 2.3, which is an application of PL-Morse theory that allows us to compute the relative acyclicity and connectivity of certain pairs of complexes. The latter result is the main application of PL–Morse theory used throughout the paper.

2.1. Some terminology and algebraic topology facts.

Let X be a topological space with a basepoint $x \in X$ and let $Y \subseteq X$ be a subspace with $x \in Y$. Let $n \geq 0$ be a nonnegative integer. We say that X is *n–connected* if $\pi_k(X, x) = 0$ for every $k \leq n$. We say that X is *n–acyclic* if $\widetilde{H}_k(X; \mathbb{Z}) = 0$ for every $k \leq n$. We say that the pair (X, Y) is *relatively n–connected* if $\pi_k(X, Y) = 0$ for every $k \leq n$. We say that the pair (X, Y) is *relatively n–acyclic* if $H_k(X, Y) = 0$

for every $k \leq n$. The *acyclicity of* X is the maximal m such that X is m–acyclic, and similarly the *connectivity of* X is the maximal m such that X is m–connected.

Notation. If X is a topological space, we will denote the connectivity and acyclicity of X by $c(X)$ and $\mathfrak{a}(X)$ respectively.
We require the following standard fact from algebraic topology.

Lemma 2.1. *Let* $X_1, \ldots, X_n$ *be a collection of topological spaces.*

(a) If each X_i *is* k_i*–connected, then the join*
$$X_1 * \ldots * X_n$$
is $(-2 + \sum_{i=1}^n (k_i + 2))$*–connected.*

(b) If each X_i *is* k_i*–acyclic, then the join*
$$X_1 * \ldots * X_n$$
is $(-2 + \sum_{i=1}^n (k_i + 2))$*–acyclic.*

2.2. **PL–Morse theory.** We now discuss the basics of PL–Morse theory and prove Lemma 2.3.

PL-Morse functions. Let X be a simplicial complex and Y a subcomplex of X. A *PL-Morse function* on X is a function $W : X^{(0)} \to \mathbb{Z}_{\geq 0}$. We define the *min-set* of W to be $M(W) = W^{-1}(\{0\})$.

Remark. In Bestvina's formulation of PL-Morse theory, it is assumed that two vertices in X with the same weight and positive weight are not adjacent. For our purposes it is not necessary to assume this, so it is not part of our definition.

Descending links. Let X be a simplicial complex equipped with a PL–Morse function W. Let $\sigma \subseteq X$ be a cell, and let $\text{lk}(\sigma)$ denote the link of σ in X. If σ is a cell of X such that W is positive and constant on the vertices of σ, then we say that σ is a W*–constant cell.* If σ is a W–constant cell of X, then the descending link $d_W(\sigma)$ is the subcomplex of $\text{lk}(\sigma)$ generated by vertices $w \in X$ such that $W(w) < W(v)$ for all vertices v of σ. Similarly, the simplicial star of σ will be denoted by $\text{st}(\sigma)$. The descending star $s_W(\sigma)$ is the join $\sigma * d_W(\sigma)$.

Connectivity and acyclicity of PL-Morse functions. Let W be a PL-Morse function on a simplicial complex X. Suppose there is a positive integer n such that for every positive weight W–constant k-cell σ, the descending link $d_W(\sigma)$ is $(n - k)$–connected. In this case, we will say that W is an n*–connected PL-Morse function.* Similarly, if there is a positive integer n such that for W–constant k–cell σ, the descending link $d_W(\sigma)$ is $(n - k)$–acyclic, we say that W is an n*–acyclic PL-Morse function.*

We have the following general result about PL–Morse functions.

Lemma 2.2. *Let X be a finite–dimensional, countable simplicial complex equipped with a PL-Morse function W. Let $Y = M(W)$.*

(a) If W is n–connected, then the pair (X, Y) is relatively $(n+1)$–connected.
(b) If W is n–acyclic, then the pair (X, Y) is relatively $(n+1)$–acyclic.

Proof. We will begin by constructing a double–indexed filtration of X. If σ is a k-cell, we denote by $W(\sigma)$ the maximum value that W obtains on σ. For $k \geq 0$ and $m > 0$, we set

$$X_{k,m} = W^{-1}([0,m)) \cup \{\sigma : W(\sigma) \leq m, \dim(\sigma) \leq k\}.$$

We will use the notation $W_{\infty,m}$ to mean the full subcomplex of X generated by vertices of weight $\leq m$. Our filtration is ordered by $X_{k,m} \subseteq X_{k',m'}$ if either $m' > m$ or $m' = m$ and $k' \geq k$. We will prove part (a) of the lemma. The proof of (b) follows from a similar argument.

The proof of (a). Since the filtration

$$X_{0,0} \subseteq X_{0,1} \subseteq \dots$$

is well-founded, it suffices to prove that for every $k \leq n+1$, the following hold:

(1) The pair $(X_{k,m}, X_{k-1,m})$ is relatively $(n+1)$–connected for every $k > 0$ and $m > 0$.
(2) The pair $(X_{0,m}, X_{\infty,m-1})$ is relatively $(n+1)$–connected for every $m > 0$.

Since (1) and (2) follow by similar reasoning, we only prove (1).

The proof of (1). Fix a $k > 0$ and $m > 0$. Let $\mathcal{T}^{k,m}$ be the set of W–constant k-cells of weight m. The complex $X_{k,m}$ is constructed from $X_{k-1,m}$ by attaching, for each $\sigma \in \mathcal{T}^{k,m}$, the complex $\sigma * d_W(\sigma)$ to $\partial\sigma * d_W(\sigma)$. By hypothesis, $d_W(\sigma)$ is $(n-k)$–connected. By Lemma 2.1, the join $\partial\sigma * d_W(\sigma)$ is $((k-2)+(n-k)+2)$–connected and hence n–connected. There is a countable filtration of $X_{k,m}$ given by arbitrarily indexing the cells of $\mathcal{T}^{k,m}$ by the natural numbers $\mathbb{N}$, and then attaching cells $\tau \in \mathcal{T}^{k,m}$ in this order to $X_{k-1,m}$. We will notate the jth term of this filtration by $X^j_{k-1,m}$. If τ is the cell added to go from $X^{j-1}_{k-1,m}$ to $X^j_{k-1,m}$, then the pair

$$\left(X^{j-1}_{k-1,m} \sqcup_{\partial\tau * d_W(\tau)} \tau * d_W(\tau), X^{j-1}_{k-1,m}\right)$$

is relatively $(n+1)$–connected. Therefore $(X^j_{k-1,m}, X^{j-1}_{k-1,m})$ is relatively $(n+1)$–connected for every $j \geq 1$, so the pair $(X_{k,m}, X_{k-1,m})$ is relatively $(n+1)$–connected. □

We have the following consequence of Lemma 2.2.

Lemma 2.3. *Let X be a finite–dimensional, countable simplicial complex with a PL-Morse function W and let $Y = M(W)$. Let n be a nonnegative integer.*

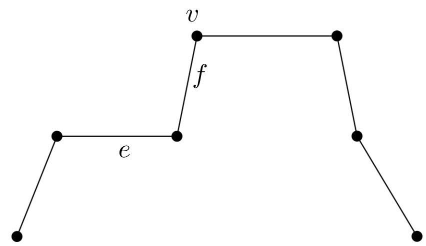

Figure 1. The filtration $X_{k,m}$. The edge e is in $X^{1,1}$, the vertex v is in $X^{0,2}$, the edge f is in $X_{0,2}$.

(a) If W is n–connected and Y is $(n+1)$–connected, then X is $(n+1)$–connected.
(b) If W is n–connected and X is n–connected, then Y is n–connected.
(c) If W is n–acyclic and Y is $(n+1)$–acyclic, then X is $(n+1)$–acyclic.
(d) If W is n–acyclic and X is n–acyclic, then Y is n–acyclic.

Proof of Lemma 2.3. We prove (a). The others follows by a similar argument. It is enough to prove that for $i \leq n$, we have $\pi_i(Y) = 0$. Fix some $i \leq n$. We have the long exact sequence in relative homotopy associated to the pair (X, Y), a part of which is given by

$$\ldots \to \pi_{i+1}(X, Y) \to \pi_i(Y) \to \pi_i(X) \to \ldots$$

Then $\pi_{i+1}(X, Y) = 0$ by Lemma 2.2 and $\pi_i(X) = 0$ by hypothesis, so $\pi_i(Y) = 0$. □

3. Covers indexed by simplicial complexes

The main output of this section is Proposition 3.1, which is a result about the acyclicity of simplicial complexes X covered by simplicial subcomplexes that are indexed by the cells of another simplicial complex Y. Proposition 3.1 is a homological version of a result of Mirzaii and van der Kallen [MvdK02]. Proposition 3.1 can also be viewed as a specific application of the homotopy colimit spectral sequence [Dug].

Bi–cellular covers. Let X and Y be two simplicial complexes. Let $\mathcal{U}(X)$ be the set of all subcomplexes of X, and let $L_X : Y^{(0)} \to \mathcal{U}$ be a function. If σ is a k-cell of Y with vertices $y_0, \ldots, y_k$, we denote

$$L_X(\sigma) = L_X(y_0) \cap \ldots \cap L_X(y_k).$$

If the set $\{L_X(y)\}_{y \in Y^{(0)}}$ covers X, we will call such an L_X a *Y-indexed cover of X*. If L_X is a Y-indexed cover of X, then there is an associated X-indexed cover of Y given by setting $L_Y(x)$ to be the full subcomplex of Y generated by vertices y such that $x \in L_X(y)$ for x a vertex of X. We say that a Y-indexed cover of X is *cellularly n–acyclic* if for each k-cell σ of Y, the complex

$L_X(\sigma)$ is $(n-k)$–acyclic. We say L_X is *bi-cellularly n–acyclic* if L_X and L_Y are both cellularly n–acyclic.

Example. We will discuss the typical way in which this notion appears in the paper. Let A be a simplicial complex, and let $X, Y \subseteq A$ be two subcomplexes. Let $\mathcal{U}(X)$ be the set of subcomplexes of X. There is a function $L_X : Y^{(0)} \to \mathcal{U}(X)$ given by $L_X(y) = \text{st}(y) \cap X$, and similarly there is a function L_Y. In our applications, we will choose A and $X, Y \subseteq A$ such that L_X and L_Y are Y–indexed and X–indexed covers respectively.

The main goal of Section 3 is to prove the following result.

Proposition 3.1. *Let X and Y be finite–dimensional, countable simplicial complexes. Let L_X be a Y-indexed cover of X. Suppose there is an integer n such that the following hold:*

- *L_X is bi-cellularly n–acyclic and*
- *Y is n–acyclic.*

Then X is n–acyclic.

We will introduce the main object that we use to prove Proposition 3.1.

Bi–cellular spectral sequence. Let X and Y be simplicial complexes, and let L_X be a Y–indexed cover of X. We have a bi-graded double complex called the *bi-cellular complex* given by

$$C_{p,q} = \bigoplus_{\sigma \in Y^{(p)}, \tau \in X^{(q)} : \tau \in L_X(\sigma)} \mathbb{Z}\sigma \times \tau.$$

The downward differential is given by the differential in X, and the leftward differential is given by the differential in Y. Let $\mathbb{E}^{*,\leftarrow}_{*,*}$ and $\mathbb{E}^{*,\downarrow}_{*,*}$ denote the leftward and downward spectral sequences associated to $C_{*,*}$. We refer to these respectively as the *leftward* and *downward bi-cellular spectral sequences.*

The leftward versus downward strategy. We now briefly review a standard application of spectral sequences, used for example to show that the G–equivariant homology of a contractible CW–complex X converges to the group homology of G when G acts on X without rotations [Bro94, Section VII]. Let $C_{p,q}$ be a double complex. There are two spectral sequences associated to $C_{p,q}$, which are the leftward and downward spectral sequences. We denote these $\mathbb{E}^{*,\leftarrow}_{*,*}$ and $\mathbb{E}^{*,\downarrow}_{*,*}$. These two spectral sequences are each constructed out of $C_{p,q}$ by two different filtrations of the total complex C_{p+q}. The key point is that $\mathbb{E}^{*,\leftarrow}_{*,*}$ and $\mathbb{E}^{*,\downarrow}_{*,*}$ both converge to filtrations of the total homology of $C_{p,q}$. Denote the total homology by $H_{p+q}(C_{*,*})$. Suppose that we want to compute some of the groups $\mathbb{E}^{2,\leftarrow}_{p,q}$. The strategy is as follows:

(1) Show that $\mathbb{E}^{*,\downarrow}_{p,q}$ converges to 0 in a range $0 < p+q < n$. This implies that $H_{p+q}(C_{*,*})$ converges to 0 for $0 < p+q < n$.
(2) Use the fact that $\mathbb{E}^{*,\leftarrow}_{p,q}$ must also converge to 0 for $0 < p+q < n$ to say something about the groups $\mathbb{E}^{2,\downarrow}_{p,q}$.

Proof of Proposition 3.1. We will apply the leftward versus downward strategy. In particular, we will show that the downward bi-cellular spectral sequence converges to $\mathbb{Z}$ for $p+q=0$ and converges to 0 for $0 < p+q \leq n$, and that the leftward bi-cellular spectral sequence converges to $H_{p+q}(X;\mathbb{Z})$ for $0 \leq p+q \leq n$. This completes the proof since the leftward and downward sequence both converge to the total homology of $C_{*,*}$.

The downward sequence. On page 1 of $\mathbb{E}^{*,\downarrow}_{*,*}$, we have

$$\mathbb{E}^{1,\downarrow}_{p,q} = \bigoplus_{\sigma \in Y^{(p)}} H_q(L_X(\sigma);\mathbb{Z}).$$

By assumption $L_X(\sigma)$ is $(n-p)$–acyclic for any p–cell σ, so for $0 \leq p+q \leq n$ and $q > 0$ we have

$$\mathbb{E}^{1,\downarrow}_{p,q} = 0.$$

Then for $q = 0$ in this range we have

$$\mathbb{E}^{1,\downarrow}_{p,*} = C_p(Y).$$

Hence $\mathbb{E}^{2,\downarrow}_{p,q}$ is $H_{p+q}(Y;\mathbb{Z})$ for $0 \leq p+q \leq n$, which is $\mathbb{Z}$ when $p+q=0$ and 0 for $0 < p+q \leq n$.

The leftward sequence. By construction, there is a canonical isomorphism

$$C_{p,q} \cong \bigoplus_{\sigma \in X^{(q)}} C_p(L_Y(\sigma)).$$

Then by the same argument as the downward case, the sequence $\mathbb{E}^{*,\leftarrow}_{*,*}$ converges to $H_{p+q}(X;\mathbb{Z})$ for $0 \leq p+q \leq n$. Since $\mathbb{E}^{*,\leftarrow}_{p,q}$ and $\mathbb{E}^{*,\downarrow}_{p,q}$ both converge to the total homology of $C_{p,q}$, we have $H_0(X;\mathbb{Z}) \cong \mathbb{Z}$ and $H_{p+q}(X;\mathbb{Z}) = 0$ for $0 < p+q < n$. □

4. Brendle–Broaddus–Putman flow

Our goal in this section is to prove Proposition 4.1, which is a technical result about how adding boundary components to a surface increases the acyclicity of certain subcomplexes of the curve complex. Let $S = S^b_g$ be a surface with $b \geq 2$. Let p_0 be a boundary component of S. Let S' be a surface and let $\iota : S \to S'$ be an embedding such that $S' \setminus \iota(S)$ is a disk and such that $\iota(p_0)$ bounds a disk in S'. Let $\mathcal{C}(S)$ denote the curve complex of S. Let $\mathcal{C}'(S)$ be the subcomplex of $\mathcal{C}(S)$ generated by curves δ such $\iota(\delta)$ is essential. Let $K(S) \subseteq \mathcal{C}(S)$ be the subcomplex generated by vertices δ such that $\iota(\delta)$ is inessential. There is a pushforward map

$$\iota_* : \mathcal{C}'(S) \to \mathcal{C}(S').$$

Brendle, Broaddus and Putman [BBM10] use a flow argument to show that the map ι_* is a homotopy equivalence. This is a similar technique to methods used by Hatcher [Hat91] and Bell and Margalit [BM06]. We describe their technique here.

Hatcher flow. Associated to each vertex $\delta \in K(S)$, there is an oriented arc α connecting a boundary component p_1 to p_0 as in Figure 2. Let $\mathrm{lk}(\delta)$ denote the link of δ in $\mathcal{C}(S)$. If $\delta \in K$,

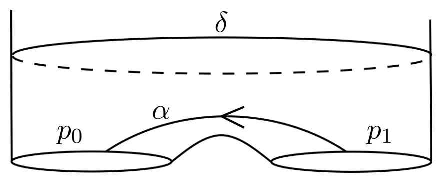

FIGURE 2. The arc α associated to the curve $\delta \in K$. α is oriented from p_1 to p_0.

then $\mathrm{lk}(\delta) \subseteq \mathcal{C}'(S)$. Then $\mathcal{C}'(S)$ is homotopy equivalent to $\mathrm{lk}(\delta)$. The homotopy equivalence is given iteratively surgering β along the arc α. An example of this surgery is given in Figure 3.

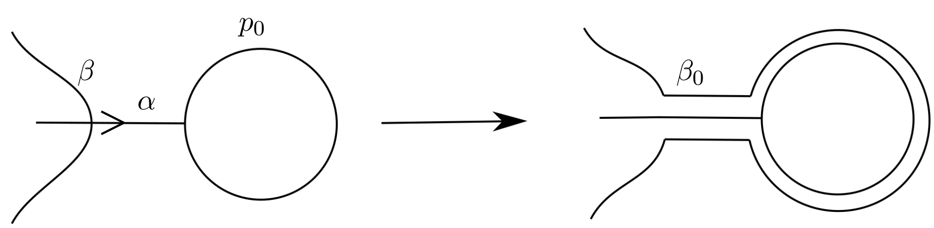

FIGURE 3. Surgering on boundary components [BBP23].

Brendle, Broaddus and Putman show that there is a homotopy equivalence

$$K(S) * \mathcal{C}(S') \simeq \mathcal{C}(S).$$

We will say that a complex $X \subseteq C(S)$ is Brendle–Broaddus–Putman compatible with a boundary component $p \in \pi_0(\partial S)$ if for every $\delta \in K(S) \cap X$ and every $\beta \in X^{(0)} \setminus (K(S) \cap X)$, the surgery of β along an arc α as in Figure 2 is also in X.

Remark. This is not *a priori* well–defined. We will more precisely state the surgery construction and prove that it is well defined in Lemma 4.2.

The main goal of the section is to prove the following result.

Proposition 4.1. *Let $S = S_g^b$ with $b \geq 2$ and $\chi(S) \leq -1$. Let p_0 be a boundary component of X. Let $X \subseteq \mathcal{C}(S)$ be Brendle–Broaddus–Putman compatible relative to p_0. Let $\iota : S \to S'$ be the inclusion map from S to the surface S' with p_0 filled in by a disk. Let $\mathcal{C}'(S)$ denote the subcomplex of the curve complex generated by curves δ such that δ remains essential when p_0 is filled in with a disk, and let $K(S)$ denote the remaining vertices of $\mathcal{C}'(S)$. Then*

(a) For every $\delta \in K(S)$, the inclusion $X \cap \mathcal{C}'(S) \hookrightarrow \mathrm{lk}(\delta)$ is a homotopy equivalence.
(b) $(X \cap K(S)) * (X \cap \mathcal{C}'(S)) \simeq X$.

Before we can prove this proposition, we will give some precise definitions.

Surgery along arcs. We now explicitly construct the surgery of β along δ. Let $S = S_g^b$ be a compact oriented surface with $b \geq 2$. Let p_0, p_1 be two boundary components of S, and let $\mathcal{C}(S)$, $\mathcal{C}'(S)$ and $K(S)$ be as above. Let $\delta \in K(S)$ be a curve and let $\beta \in \mathcal{C}'(S)$ be another curve. We will define the *surgery of β along δ*, which we denote $\mathrm{surg}_\delta(\beta)$. Let $\widehat{\delta}$ be a smooth representative of δ, and let $\widehat{\alpha}$ be an oriented arc connecting p_1 to p_0 such that α is disjoint from $\widehat{\delta}$. Choose a representative $\widehat{\beta}$ of β such that the intersection number $\left|\widehat{\beta} \cap \widehat{\alpha}\right|$ is minimal. If $\widehat{\beta}$ is disjoint from $\widehat{\alpha}$, then we say $\mathrm{Surg}_\delta(\beta) = \beta$. Otherwise, let q be the point of intersection of $\widehat{\beta}$ with $\widehat{\alpha}$ that is closest along $\widehat{\alpha}$ to the boundary component p_0. Let U be an annulus with $p_0 \subseteq \partial U$. Let $\widehat{\alpha}_{q,p_0}$ denote the embedded interval given by restricting $\widehat{\alpha}$ to the sub–interval connecting q to p_0. Let V be an embedded copy of $I \times I$ such that $I \times \frac{1}{2} = \alpha_{q,p_0}$. Define

$$\mathrm{Surg}_\delta(\beta) = \text{ the isotopy class of } \widehat{\beta}\Delta(\partial(I \cup U))$$

where Δ denotes symmetric difference. We have the following result.

Lemma 4.2. *Let $S, S', p_0, p_1, \delta, \beta$ be as above. The curve $\mathrm{Surg}_\delta(\beta)$ is well–defined.*

Proof. In the construction of $\mathrm{Surg}_\delta(\beta)$, there are six choices made, namely $\widehat{\delta}$, $\widehat{\alpha}$, $\widehat{\beta}$, U, and I. Only the choice of $\widehat{\alpha}$ is still a choice up to isotopy, since two arcs $\widehat{\alpha}$ and $\widehat{\alpha}'$ need not be isotopic. Since we are working up to isotopy, we may assume that $\widehat{\alpha}$ and $\widehat{\alpha}'$ have the same endpoints. Let α and α' denote the isotopy classes of the arcs $\widehat{\alpha}$ and $\widehat{\alpha'}$. There are integers $m, n \in \mathbb{Z}$ such that $T_{p_0}^m T_{p_1}^n \alpha = \alpha'$, where T_{p_i} denotes the Dehn twist along a curve cobounding an annulus with p_i. But then if γ denotes the surgery of β along α up to isotopy and γ' for β along α' up to isotopy, we see that $T_{p_0}^m T_{p_1}^n \gamma = \gamma'$. But T_{p_0} and T_{p_1} act trivially on $\mathcal{C}(S)$, so $\mathrm{Surg}_\delta(\beta)$ is well–defined. □

Brendle–Broaddus–Putman compatibility. Let $S = S_g^b$ with $b \geq 2$. Let $X \subseteq \mathcal{C}(S)$ be a subcomplex of the curve complex. Let p_0 be a boundary component of S. Let S', $\mathcal{C}'(S)$ and $K(S)$

be as above. We say that X is *Brendle–Broaddus–Putman compatible relative to* p_0 if for every curve $\delta \in X \cap K(S)$ and every $\beta \in X \cap \mathcal{C}'(S)$, we have $\mathrm{Surg}_\delta(\beta) \in X$. We say that X is *Brendle–Broaddus–Putman compatible* if it is Brendle–Broaddus–Putman compatible relative to every boundary component of S.

Before proving Proposition 4.1, we require the following lemma from Brendle, Broaddus and Putman [BBP23, Lemma 4.1].

Lemma 4.3. *Let X be a simplicial complex, let I be a discrete set, and let $Y \subseteq I * X$ be a subcomplex such that $I, X \subseteq Y$. Suppose that, for any $i \in I$, the inclusion $\mathrm{lk}_Y(i) \hookrightarrow X$ is a homotopy equivalence. Then the inclusion $Y \hookrightarrow I * X$ is a homotopy equivalence.*

We are now ready to conclude Section 4.1. Our proof of Proposition 4.1 is similar to the argument of Brendle, Broaddus and Putman [BBP23].

Proof of Proposition 4.1. Part *(b)* of the lemma follows from part *(a)* and Lemma 4.3, so it suffices to prove part (a). Let $\delta \in K(S)$. Since $X \cap \mathcal{C}'(S)$ and $X \cap \mathrm{lk}(\delta)$ both have the homotopy type of CW–complexes, it suffices to show that the inclusion map $\kappa : X \cap \mathrm{lk}(\delta) \to X \cap \mathcal{C}'(S)$ induces an isomorphism on homotopy groups. Equivalently, it suffices to show that the relative homotopy groups

$$\pi_k(X \cap \mathcal{C}'(S), X \cap \mathrm{lk}(\delta))$$

vanish for all $k \geq 0$.

Choose a hyperbolic metric on S. Then for each $\beta \in X \cap \mathcal{C}'(S)$, let $\widehat{\beta}$ be the geodesic representative. Let $\widehat{\alpha}$ be an oriented geodesic arc connecting p_1 to p_0 and disjoint from δ. Since these are geodesics representatives, any pair of chosen representatives intersects minimally [FM12]. By perturbing the $\widehat{\beta}$ slightly, we may additionally assume that:

- each $\widehat{\beta}$ intersects $\widehat{\alpha}$ in a different point than any other $\widehat{\beta}$, and
- $\widehat{\beta}$ minimally intersects $\widehat{\beta}'$ for any $\beta, \beta' \in \mathcal{C}(S)$.

Let $W(\beta)$ denote the distance along α between p_1 and q, where q is the closest point of intersection of β and α with p_0. Let $\psi : S^k \to X \cap \mathcal{C}'(S)$ be a simplicial representative of a class in $\pi_k(X \cap \mathcal{C}'(S), X \cap \mathrm{lk}(\delta))$, where S^k is a simplicial decomposition of a k–sphere. Let $\beta \in \psi(S^k)$ be a vertex with $W(\beta)$ maximal. If $W(\beta) = 0$ then we are done, so assume otherwise. Note that $d_W(\beta)$ is a cone with cone point $\mathrm{Surg}_\delta(\beta)$. Hence ψ is homotopic to a map $\psi' : S^k \to X \cap \mathcal{C}'(S)$ such that the maximal intersection number of any vertex in $\psi'(S^k)$ is not larger than for $\psi(S^k)$. Since iteratively applying Surg_δ to a curve β eventually stabilizes in a curve $\overline{\beta}$ with $W(\overline{\beta})) = 0$, the above process eventually terminates in ψ with both $\psi \simeq \psi''$ and $\psi'' : S^k \to X \cap \mathrm{lk}(\delta)$, so the lemma holds. ∎

5. The Complex of splitting curves

Our main goal in this section is to prove Proposition 5.1, which says that the complex of splitting curves is highly acyclic. We begin with a pair of posets of surfaces equipped with certain decorations on the boundary components, and then we will state Proposition 5.1. In Section 5.1, we will prove Lemma 5.2, which is an acyclicity result about a complex called the nonseparating arc complex. We will conclude the section with Section 5.2, where we will prove Proposition 5.1.

Cutting curves on surfaces. Let S be a surface and $S' \subseteq S$ a compact submanifold. The notation $S \searrow S'$ denotes Farb and Margalit's notion of cutting subsurfaces [FM12].

Partitioned surface. Following Putman [Put07], a *partitioned surface* $\Sigma = (S, P)$ is a pair consisting of a compact, connected, oriented surface S and a partition P of the set of boundary components of S. A *block of a partition* is one set in the partition. There is a poset $\mathfrak{TSurPos}$ of partitioned surfaces where $(S, P) \leq (S', P')$ if there is an embedding $\iota : S \to S'$ such that:

- for each block $B \in P$, there is $S_B \in \pi_0(S' \searrow \iota(S))$ with $\iota(B) \subseteq S_B$, and
- for each $\widehat{S} \in \pi_0(S' \searrow \iota(S))$, the intersection $\partial\widehat{S} \cap \iota(S)$ is a block of P.

Complex of separating curves. The *complex of separating curves* $\mathcal{C}_{\text{sep}}(\Sigma)$ is the full subcomplex of $\mathcal{C}(S)$ generated by separating curves δ such that each block $B \in P$ is contained entirely in one connected component of $S \searrow \delta$.

Vertex complement. Following Hatcher and Margalit [HM12], a *vertex complement* is a partitioned surface $\Sigma = (S, P)$ such that:

- $|P| = 2$, and
- the two blocks of P come equipped with labels. One is labeled B_+ and the other is labeled B_-.

We will denote a vertex complement by

$$(S, P, B_+, B_-).$$

The terminology "vertex complement" refers to Σ being the complement of a vertex of a certain complex called the complex of minimizing cycles, which we discuss in Section 7. We will say that Σ is a *vertex complement on a surface* S if the underlying surface of Σ is S. The set of vertex complements is a poset which we will denote $\mathfrak{VerCom}$. We have $\Sigma \leq \Sigma'$ if:

- $\Sigma \leq \Sigma'$ in $\mathfrak{TSurPos}$, and

- there is an inclusion $\iota : S \to S'$ that realizes $\Sigma \leq \Sigma'$ in $\mathfrak{TSurPos}$ such that the two connected components of $S' \setminus \iota(S)$ are S_+ and S_-, with $\partial S_+ = B_+ \cup B'_+$ and $\partial S_- = B_- \cup B'_-$.

We will denote the genus of the underlying surface S in $\Sigma = (S, P, B_+, B_-)$ by $g(\Sigma)$.

The complex of splitting curves. Let $\Sigma = (S_g^b, P, B_+, B_-)$ be a vertex complement. The *complex of splitting curves* $\mathcal{C}_{\text{split}}(S, P)$ is the full subcomplex of $\mathcal{C}_{\text{sep}}(S, P)$ generated by curves δ such that B_+ and B_- are contained in separate connected components of $S \setminus \delta$.

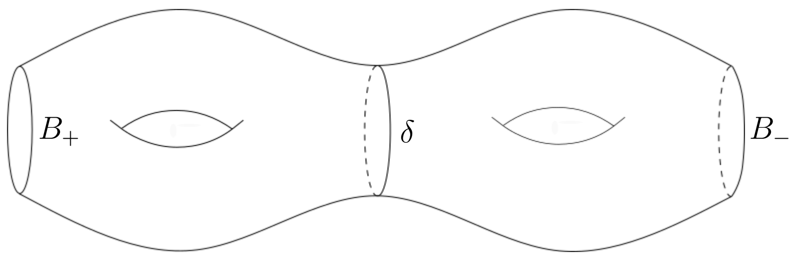

FIGURE 4. A splitting curve.

The remainder of this section will be devoted to the proof of the following proposition.

Proposition 5.1. *Let $\Sigma = (S_g^b, P, B_+, B_-)$ be a vertex complement. The acyclicity $\mathfrak{a}(\mathcal{C}_{\text{split}}(\Sigma))$ satisfies*

$$\mathfrak{a}(\mathcal{C}_{\text{split}}(\Sigma)) \geq g - 3 + \mathbb{1}_{|B_+| \geq 2} + \mathbb{1}_{|B_-| \geq 2}.$$

We began by proving a result about a variant of the arc complex called the nonseparating arc complex. We complete the proof of Proposition 5.1 via induction and a spectral sequence argument based on the content of Section 3.

5.1. **The nonseparating arc complex.** The goal of this section is to prove Lemma 5.2, which is a result about the acyclicity of a certain subcomplex of the arc complex called the nonseparating arc complex. We begin by defining a poset of compact surfaces with marked compact intervals on the boundary of the surface. We then prove Lemma 5.3, which is an auxiliary result about the contractibility of another subcomplex of the arc complex. We conclude by leveraging Lemma 5.3 and Proposition 3.1 to prove Lemma 5.2.

Marked vertex complements. The poset $\mathfrak{MarkSur}$ has elements consisting of vertex complements Σ equipped with a set of compact intervals Q on the boundary components in B_+ such that:

- $|Q| \geq 2$, and
- each $p \in B_+$ contains an interval of Q.

If $p \in B_+$ is a boundary component, we let $Q(p)$ denote the set of intervals contained in p. We say that $(\Sigma, Q) \leq (\Sigma', Q')$ if $\Sigma \leq \Sigma'$ in $\mathfrak{VerCom}$.

The nonseparating arc complex. Let $(\Sigma, Q) \in \mathfrak{MarkSur}$. The *nonseparating arc complex* $\mathcal{A}_{\text{nosep}}(\Sigma, Q)$ is the subcomplex of the arc complex [Sch15] of S consisting of cells σ such that:

(1) $S \searrow \sigma$ is connected, and
(2) every arc of σ has endpoints in distinct intervals of Q.

Note that the arc complex is a complex of arcs up to isotopy. We will allow isotopies that move endpoints as long as the endpoints of the arcs remain in the intervals in Q. The goal of Section 5.1 is to prove the following lemma.

Lemma 5.2. *Let $S = S_g^b$ be a compact surface and let $(\Sigma, Q) \in \mathfrak{MarkSur}$ with underlying surface S. The complex $\mathcal{A}_{\text{nosep}}(\Sigma, Q)$ is $(g + |B_+| - 3)$–acyclic.*

The proof proceeds by applying Proposition 3.1 to the complex $\mathcal{A}_{\text{nosep}}(\Sigma, Q)$ and the *complex of Q–arcs* $\mathcal{A}(\Sigma, Q)$, which is the complex of arcs with endpoints in distinct intervals in Q. For $\mathcal{A}(\Sigma, Q)$, we will allow isotopies of arcs to change the endpoints of α, but the endpoints must remain inside the intervals of Q throughout the isotopy.

The outline of Section 5.1. We will first discuss the process of cutting open an arc α in $(\Sigma, Q) \in \mathfrak{VerCom}$. We will use the method of "Hatcher flow" [Hat91] to prove Lemma 5.3 that $\mathcal{A}(\Sigma, Q)$ is contractible, and then use Proposition 3.1 to prove Lemma 5.2.

Cutting open arcs in (Σ, Q). Let $(\Sigma, Q) \in \mathfrak{MarkSur}$ and let $\alpha \in \mathcal{A}(\Sigma, Q)$ be a vertex. We will define $(\Sigma, Q) \searrow \alpha \in \mathfrak{MarkSur}$. Let $\Sigma = (S, P, B_+, B_-)$. The object $\Sigma \searrow \alpha \in \mathfrak{VerCom}$ is given by the unique vertex complement structure on $S \searrow \alpha$ such that:

- $\Sigma \searrow \alpha < \Sigma$, and
- the inclusion $\iota : \Sigma \searrow \alpha \to \Sigma$ that realizes $\Sigma \searrow \alpha < \Sigma$ is the inclusion $S \searrow \alpha \to S$.

We will denote $\Sigma \searrow \alpha = (S \searrow \alpha, P^\alpha, B_+^\alpha, B_-^\alpha)$. Let $(\Sigma \searrow \alpha, Q_\alpha) = (\Sigma, Q) \searrow \alpha$. Q_α is defined as follows. Let $\iota : S \searrow \alpha \to S$ denote the natural inclusion. If $p \in B_+^\alpha$ is a boundary component with $\iota(p)$ isotopic to a boundary component $p' \in B_+$, then we will define Q_α to have $|Q_\alpha(p)| = |Q_\alpha(p')|$. If $\iota(p)$ is not isotopic to a boundary component of S, we have two cases.

- Case 1: there is an embedded $S_{0,3} \hookrightarrow S$ with $\partial S_{0,3} = \iota(p) \cup p_0 \cup p_1$ with $p_0, p_1 \subseteq \partial S$. In this case, α has endpoints in p_0 and p_1. Let $f : p \to p_0 \cup \alpha \cup p_1$ be a smooth map that extends to a smooth map $F : S \searrow \alpha \to S$ which is homotopic to ι. We define

$$Q_\alpha(p) = \pi_0 \left(f^{-1}(\alpha) \bigcup_{I \in Q(p_0) \cup Q(p_1)} f^{-1}(I) \right).$$

- Case 2: there is no such $S_{0,3} \subseteq S$. In this case, the endpoints of α are in the same boundary component $q \in B_+$. There is a unique boundary component $p' \in B_+^\alpha$ such that $\iota(p'), \iota(p)$ and q cobound an embedded $S_{0,3} \subseteq S$. Let $f : p \cup p' \to q \cup \alpha$ be a smooth map that extends to a smooth map $F : S \searrow \alpha \to S$ such that F is homotopic to ι. We define
$$Q_\alpha(p) = \pi_0 \left(\bigcup_{I \in Q(q)} f^{-1}(I) \right).$$

The motivation behind this definition is that for σ a cell of $\mathcal{A}_{\text{nosep}}(\Sigma, Q)$, we have
$$\mathrm{lk}_{\mathcal{A}_{\text{nosep}}(\Sigma,Q)}(\sigma) = \mathcal{A}_{\text{nosep}}((\Sigma, Q) \searrow \sigma)$$
and similarly for $\sigma \subseteq \mathcal{A}(\Sigma, Q)$.

We will prove the following lemma about the complex of Q–arcs $\mathcal{A}(\Sigma, Q)$. This proof follows a method of Hatcher [Hat91], which is sometimes referred to as Hatcher flow. This strategy is also used in the proof of Lemma 4.1.

Lemma 5.3. *Let $(\Sigma, Q) \in \mathfrak{MarkSur}$. The complex $\mathcal{A}(\Sigma, Q)$ is contractible.*

Proof. Fix an arc $\beta \in \mathcal{A}(\Sigma, Q)$. Let $q_0, q_1 \in Q$ be the endpoints of β and orient β from q_0 to q_1. We will define a surgery function Surg_β on the complex $\mathcal{A}(\Sigma, Q)$.

We begin by defining $\mathrm{Surg}_\beta(\alpha)$ for α a vertex of $\mathcal{A}(\Sigma, Q)$. Assume that α is isotoped to be in minimal position with β. If α is disjoint from β, then $\mathrm{Surg}_\alpha(\beta) = \alpha$. Otherwise, suppose that p is the point of intersection of α and β closest along β to the point q_1. Cut α at the point p and paste in the subsegment of β connecting p to q_1, and then homotope slightly to get 2 arcs α' and α'' with α' disjoint from α and β, and α'' disjoint from α and minimally intersecting β. Refer to Figure 5 for a picture of the construction.

Now, one of α' or α'' must have endpoints in distinct intervals in Q since both α and β both have endpoints in distinct intervals of Q. Hence, the function Surg_β is defined on vertices by:
$$\mathrm{Surg}_\beta(\alpha) = \begin{cases} \alpha' & \text{if } \alpha'' \text{ is not a vertex of } \mathcal{A}(\Sigma, Q) \\ \alpha'' & \text{if } \alpha' \text{ is not a vertex of } \mathcal{A}(\Sigma, Q) \\ \frac{1}{2}\alpha' + \frac{1}{2}\alpha'' & \text{otherwise} \end{cases}$$

We now use the function Surg_β to that the inclusion $\mathrm{st}(\beta) \to \mathcal{A}(\Sigma, Q)$ is a weak homotopy equivalence, and hence a homotopy equivalence. Since $\mathrm{st}(\beta)$ is contractible, this completes the proof of the Lemma. We use essentially the same strategy as in the proof of Lemma 4.3. Choose geodesic representatives for every arc, and perturb them so that there is a total order on closest intersection point in $\alpha \cap \beta$ to q_1 along β. Then there is a weight function given by the distance

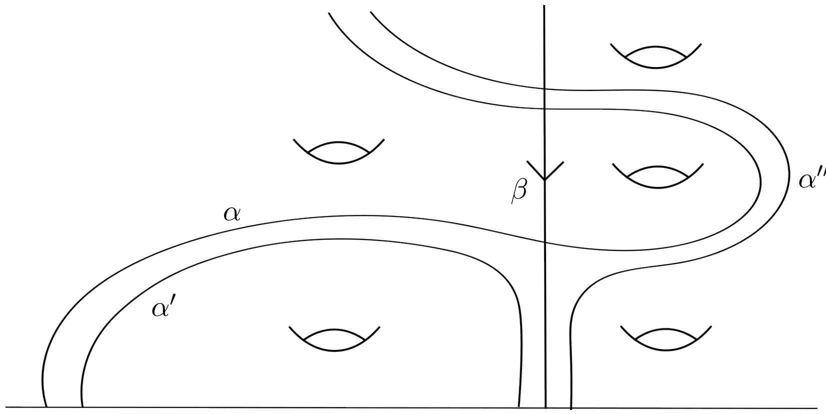

FIGURE 5. Surgery on arcs.

along β from this point of intersection to q_1. Then any map $\varphi : S^n \to \mathcal{A}(\Sigma, Q)$ can be iteratively surgered to $\psi : \mathrm{st}(\beta) \to \mathcal{A}(\Sigma, Q)$ with $\varphi \simeq \psi$. □

We are now almost ready to prove Lemma 5.2, which we recall says that $\mathcal{A}_{\mathrm{nosep}}(\Sigma, Q)$ is $(g + |B_+| - 3)$–acyclic. We first must extend the notion of cutting open curves and arcs on surfaces to the poset $\mathfrak{MarkSur}$.

Proof of Lemma 5.2. We will induct on the poset $\mathfrak{MarkSur}$.

Base cases. The base cases are given elements $(\Sigma, Q) \in \mathfrak{MarkSur}$ such that $g + |B_+| - 3 \leq -1$, or equivalently that $g + |B_+| \leq 2$. Since Q has at least two intervals, $|B_+|$ is nonempty. Hence the two cases to consider are:

- $S = S_1^b$ with $|B_+| = 1$, and
- $S = S_0^b$ with $|B_+| = 2$.

In both cases, for any choice of Q such that $(\Sigma, Q) \in \mathfrak{MarkSur}$, the complex $\mathcal{A}_{\mathrm{nosep}}(\Sigma, Q)$ is nonempty.

Inductive step. Let $\Sigma = (S, P, B_+, B_-) \in \mathfrak{VerCom}$ and $Q \subseteq \partial\Sigma$ with $(\Sigma, Q) \in \mathfrak{MarkSur}$ such that for every $(\Sigma', Q') \in \mathfrak{MarkSur}$ with $(\Sigma', Q') < (\Sigma, Q)$, the lemma holds for (Σ', Q'). We will then show that the lemma holds for (Σ, Q) as well.

We will apply Proposition 3.1 with the complexes $\mathcal{A}_{\mathrm{nosep}}(\Sigma, Q)$ and $\mathcal{A}(\Sigma, Q)$. For any k–cell σ of $\mathcal{A}_{\mathrm{nosep}}(\Sigma, Q)$, the subcomplex $L_{\mathcal{A}(\Sigma,Q)}(\sigma)$ is contractible. Since $\mathcal{A}(\Sigma, Q)$ is contractible by Lemma 5.3, it suffices to show that for any k–cell τ of $\mathcal{A}(\Sigma, Q)$, we have

$$\mathfrak{a}(L_{\mathcal{A}_{\mathrm{nosep}}(\Sigma,Q)}(\tau)) = \dim(\mathcal{A}_{\mathrm{nosep}}(\Sigma, Q)) - k - 1$$

Observe that if α is an arc on Σ with endpoints in distinct intervals in Q, then cutting Σ along α does one of three things:

(1) decreases $|B_+|$ by 1 if α joins two distinct boundary components,
(2) decreases g by 1 and increases $|B_+|$ by 1 if α has endpoints in the same boundary component, or
(3) splits Σ into two surfaces Σ' and Σ'' with

$$g(\Sigma') + g(\Sigma'') = g(\Sigma) \text{ and } |B'_+| + |B''_+| = |B_+| + 1.$$

We will refer to arcs in the first collection as Type (1) arcs, and similarly for Type (2) and Type (3). Let $\sigma \subseteq \mathcal{A}(\Sigma, \mathbb{Q})$ be a k–cell. Let $(\Sigma_1, Q_1), \ldots, (\Sigma_m, Q_m) \in \mathfrak{MarkSur}$ be the connected components of $(\Sigma, Q) \searrow \sigma$. If $m = 1$, then $\sigma \subseteq \mathcal{A}_{\text{nosep}}(\Sigma, Q)$, so $L_{\mathcal{A}_{\text{nosep}}(\Sigma,Q)}(\sigma)$ is contractible. Otherwise, let $\tau \subseteq \sigma$ be a cell of maximal dimension with the property that $(\Sigma, Q) \searrow \tau$ is connected. Let $\tau' \subseteq \sigma$ be the unique sub–cell with $\tau * \tau' = \sigma$. Let $(\Sigma, Q) \searrow \tau = (\Sigma_\tau, Q_\tau)$, with $\Sigma_\tau = (S \searrow \tau, P^\tau, B^\tau_+, B^\tau_-)$. Note that $\dim(\tau) = k + 1 - m$. Since every arc in τ must be either Type (1) or Type (2), we have

$$g(\Sigma_\tau) + |B^\tau_+| \geq g(\Sigma) + |B_+| - (k - m + 2).$$

Then by construction, every arc of τ' is of Type (3) when restricted to (Σ_τ, Q_τ). Therefore we have

$$\sum_{j=1}^{m} g(\Sigma_j) + |B^j_+| = g(\Sigma_\tau) + |B^\tau_+| + m - 1.$$

Putting this together and applying Lemma 2.1 along with the inductive hypothesis, we see that

$$\begin{aligned}
\mathfrak{a}(L_{\mathcal{A}_{\text{nosep}}(\Sigma,Q)}(\sigma) &= -2 + \sum_{j=1}^{m} \Big(g(\Sigma_j) + |B^+_j| - 3 + 2 \Big) \\
&= -2 + \sum_{j=1}^{m} \Big(g(\Sigma_j) + |B^+_j| \Big) - m \\
&= -2 + g(\Sigma_\tau) + |B^\tau_+| + m - 1 - m \\
&= +g(\Sigma_\tau) + |B^\tau_+| - 3 \\
&\geq g(\Sigma) + |B_+| - k + m - 5.
\end{aligned}$$

We have assumed that $m \geq 2$, so this last expression is bounded below by $g(\Sigma) + |B_+| - k - 3$, so the proof is complete. □

5.2. **The proof of Proposition 5.1.** We now show that $\mathfrak{a}(\mathcal{C}_{\text{split}}(\Sigma)) \geq g-3+\mathbb{1}_{|B_+|\geq 2}+\mathbb{1}_{|B_-|\geq 2}$. We will begin by describing the outline of the proof of Proposition 5.1. We will then carry out some of the steps in the outline in a pair of lemmas. We will conclude Section 5.2, and Section 5, by proving Proposition 5.1.

The setup of the proof of Proposition 5.1. As in the proof of Lemma 5.2, the proof follows by induction on the poset $\mathfrak{VerCom}$. Let $\Sigma = (S, P, B_+, B_-)$ be a vertex complement such that Proposition 5.1 holds for all $\mathcal{T} < \Sigma$. Suppose without loss of generality that $|B_+| \geq |B_-|$. Choose a set Q of compact subintervals of $\partial\Sigma$ with $|Q|$ minimal such that $(S, Q) \in \mathfrak{MarkSur}$. If σ is a cell of $\mathcal{C}_{\text{split}}(\Sigma)$, let $L_{\mathcal{A}_{\text{nosep}}(\Sigma,Q)}(\sigma)$ be the full subcomplex of $\mathcal{A}_{\text{nosep}}(\Sigma, Q)$ generated by arcs disjoint from σ. Similarly, for for cells $\tau \in \mathcal{A}_{\text{nosep}}(\Sigma, Q)$, let $L_{\mathcal{C}}(\tau)$ denote the subcomplex of $\mathcal{C}_{\text{split}}(\Sigma)$ generated by cells σ with σ disjoint from τ. Let $C_{p,q}$ be the bi-cellular complex

$$\bigoplus_{\sigma\in\mathcal{A}_{\text{nosep}}(\Sigma,Q)^{(p)}} C_q(L_C(\sigma)).$$

Let $\mathbb{E}^{*,\downarrow}_{*,*}$ and $\mathbb{E}^{*,\leftarrow}_{*,*}$ be the downward and leftward bi-cellular spectral sequences respectively. The goal is to prove the following two facts:

(1) the downward sequence converges to $\mathbb{Z}$ for $p+q=0$ and converges to 0 for $0 < p+q < \dim(\mathcal{C}_{\text{split}}(\Sigma))$, and
(2) the leftward sequence converges to $H_{p+q}(\mathcal{C}_{\text{split}}(\Sigma))$ for $0 \leq p+q < \dim(\mathcal{C}_{\text{split}}(\Sigma))$.

Given these two facts, $\widetilde{H}_{p+q}(\mathcal{C}_{\text{split}}(\Sigma)) = 0$ for $0 \leq p+q < \dim(C_{\text{split}}(\Sigma))$. Fact (1) follows from essentially the same argument as in the proof of Proposition 3.1, so the bulk of Section 5.2 will be devoted to proving Fact (2), which we will record as the following result.

Lemma 5.4. *Let $\Sigma \in \mathfrak{VerCom}$ such that for every $\mathcal{T} < \Sigma$, Proposition 5.1 holds for$\mathcal{C}_{\text{split}}(\mathcal{T})$. Let $\mathbb{E}^{1,\leftarrow}_{p,q}$ be the leftward spectral sequence discussed above. Then*

$$\mathbb{E}^{1,\leftarrow}_{p,q} \Rightarrow H_{p+q}(\mathcal{C}_{\text{split}}(\Sigma);\mathbb{Z})$$

for $p+q < \dim(\mathcal{C}_{\text{split}}(\Sigma))$.

The idea of the proof is to use a variant of PL–Morse theory for homology in coefficient systems. In particular, we will prove the following lemma.

Lemma 5.5. *Let $\Sigma \in \mathfrak{VerCom}$ be as in Lemma 5.4. Let $p, q \geq 0$ be integers with $p > 0$ and $p+q \leq g(\Sigma)-3+\mathbb{1}_{|B_+|\geq 2}+\mathbb{1}_{|B_-|\geq 2}$. Let $\delta \in \mathcal{C}_{\text{split}}(\Sigma)$. Label the connected components of $\Sigma \setminus \delta$ by Σ_+ and Σ_- where $B_+ \subseteq \Sigma_+$ and $B_- \subseteq \Sigma_-$. Suppose that $H_p(L_{\mathcal{A}_{\text{nosep}}(\Sigma,Q)}(\delta);\mathbb{Z}) \neq 0$. Then $\mathcal{C}_{\text{split}}(\Sigma_-)$ is at least $(q-1)$–acyclic.*

Note that if $\delta \in \mathcal{C}_{\text{split}}(\Sigma, Q)$, then Lemma 5.2 does not tell us that$L_{\mathcal{A}_{\text{nosep}}(\Sigma,Q)}(\delta)$ is p–acylic. The content of Lemma 5.5 is that for these δ, the acyclicity of $C_{\text{split}}(\Sigma_-)$ is high enough to allow us to carry out something resembling PL–Morse theory on the spectral sequence $\mathbb{E}^{1,\leftarrow}_{p,q}$.

Proof of Lemma 5.5. Let $p' = g(\Sigma_+) + |B_+| - 3$. Since $H_p(L_{\mathcal{A}}(\delta); \mathbb{Z}) \neq 0$, Lemma 5.2 implies that $p' < p$. Therefore we have

$$g(\Sigma) + |B_+| - 3 - p' \geq q + 1.$$

We also have $g(\Sigma_+) + g(\Sigma_-) = g(\Sigma)$. Since $\Sigma' < \Sigma$, Proposition 5.1 holds for Σ' by hypothesis. Therefore we have

$$\begin{aligned}\mathfrak{a}(\mathcal{C}_{\text{split}}(\Sigma_-)) &\geq g(\Sigma_-) - 3 + \mathbb{1}_{|B_-|\geq 2} \\ &\geq g(\Sigma) - g(\Sigma_+) + \mathbb{1}_{|B_-|\geq 2} - 3.\end{aligned}$$

By rearranging $p' = g(\Sigma_+) + |B_+| - 3$, we have $-g(\Sigma_+) = -p' + |B_+| - 3$. Therefore we have

$$\mathfrak{a}(\mathcal{C}_{\text{split}}(\Sigma_-)) \geq g(\Sigma) - p' - 6 + |B_+| + \mathbb{1}_{|B-|\geq 2}$$

By applying $g(\Sigma) + |B_+| - 3 - p' \geq q + 1$, we have

$$\begin{aligned}\mathfrak{a}(\mathcal{C}_{\text{split}}(\Sigma_-)) &\geq (q+1) - 3 + |B_+| - \mathbb{1}_{|B_+|\geq 2} \\ &\geq q - 1\end{aligned}$$

so the lemma holds. □

We are now ready to prove Lemma 5.4.

Proof of Lemma 5.4. On page 1, the sequence $\mathbb{E}^{1,\leftarrow}_{p,q}$ is given by

$$\mathbb{E}^{1,\leftarrow}_{p,q} = \bigoplus_{\tau \in \mathcal{C}_{\text{split}}(\Sigma)^{(q)}} H_p(L_{\mathcal{A}}(\tau)).$$

Observe that in the column $p = 0$, the chain complex $\mathbb{E}^{1,\leftarrow}_{0,q}$ is identified with $C_*(\mathcal{C}_{\text{split}}(\Sigma))$, since Lemma 5.2 implies that each cell $\tau \in \mathcal{C}_{\text{split}}(\Sigma)$ has $L_{\mathcal{A}}(\tau)$ connected. Hence it is enough to show that $\mathbb{E}^{2,\leftarrow}_{p,q} = 0$ for $p > 0$ and $0 < p + q < \dim(\mathcal{C}_{\text{split}}(\Sigma))$.

Pick a pair p, q with $p > 0$ and $p + q < \dim(\mathcal{C}_{\text{split}}(\Sigma))$. We will show that $\mathbb{E}^{2,\leftarrow}_{p,q} = 0$. For a vertex $\delta \in \mathcal{C}_{\text{split}}(\Sigma)$, let $W(\delta) = g(\Sigma_+) + |B_+| - 3$, where Σ_+ is the connected component of $\Sigma \searrow \gamma$ that contains B_+. Let Σ_- be the connected component of the vertex complement on the connected component of $\Sigma \searrow \delta$ that contains B_-,. An example of δ, Σ_+ and Σ_- can be seen in Figure 6.

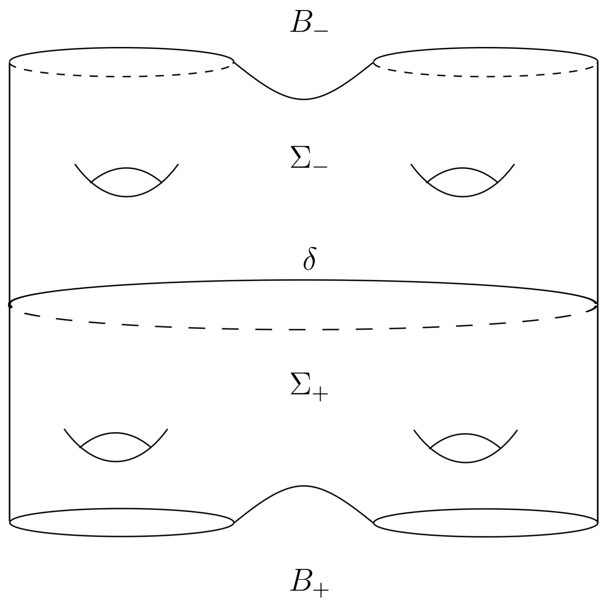

FIGURE 6. A curve δ with minimal $\dim(L_{\mathcal{A}}(\delta))$.

Let $\mathbb{E}^{1,\leftarrow}_{p,*}$ denote the chain complex in the pth column of the spectral sequence $\mathbb{E}^{1,\leftarrow}_{*,*}$. We will show that $\mathbb{E}^{2,\leftarrow}_{p,q} = H_q\left(\mathbb{E}^{1,\leftarrow}_{p,*};\mathbb{Z}\right) = 0$. There is a filtration X_k of $\mathcal{C}_{\text{split}}(\Sigma)$ given by $X_k = \bigcup_{\tau\in F_k}\tau$, where

$$F_k = \left\{\tau \subseteq \mathcal{C}_{\text{split}}(\Sigma) : \min_{\delta\in\tau^{(0)}} W(\delta) \geq p-k\right\}$$

This filtration starts at $k=0$ and runs to $k=p$. Each X_k induces a subcomplex of $\mathbb{E}^{1,\leftarrow}_{p,*}$ by

$$\mathbb{E}^{1,\leftarrow}_{p,q}(X_k) = \bigoplus_{\tau\in\mathcal{C}_{\text{split}}(\Sigma)^{(q)}:\tau\subseteq X_k} H_p(L_{\mathcal{A}}(\tau)).$$

By construction, these $\mathbb{E}^{1,\leftarrow}_{p,*}(X_k)$ are a filtration of $\mathbb{E}^{1,\leftarrow}_{p,*}$. It therefore suffices to prove the following two facts:

- $H_q(\mathbb{E}^{1,\leftarrow}_{p,*}(X_1)) = 0$, and
- $H_q(\mathbb{E}^{1,\leftarrow}_{p,*}(X_k), \mathbb{E}^{1,\leftarrow}_{p,*}(X_{k-1})) = 0$ for $1 \leq kp$.

We will prove each of these in turn.

$H_q(\mathbb{E}^{1,\leftarrow}_{p,*}(X_0);\mathbb{Z}) = 0$. Each $\tau\in X_0$ has $\mathfrak{a}(L_{\mathcal{A}}(\tau)) = p$. Hence $\mathbb{E}^{1,\leftarrow}_{p,*}(X_0)$ is identically 0.

$H_q(\mathbb{E}^{1,\leftarrow}_{p,*}(X_k), \mathbb{E}^{1,\leftarrow}_{p,*}(X_{k-1});\mathbb{Z}) = 0$ *for* $1 \leq k \leq p$. Let

$$\Delta_k = \{\delta \in \mathcal{C}_{\text{split}}(\Sigma) : W(\delta) = p-k\}.$$

Note that Δ_k is a discrete set, i.e., no two vertices in Δ_k are adjacent. Assume that Δ_k is ordered arbitrarily, which induces a filtration Y_k^j of X_k, where Y_k^j is given by attaching to X_{k-1} the first j elements in the order on Δ_k. Each Y_k^j induces a subcomplex of $\mathbb{E}^{1,\leftarrow}_{p,*}(X_k)$, which we denote $\mathbb{E}^{1,\leftarrow}_{p,*}(Y_k^j)$. We have the following claim.

Claim: For any $j \geq 1$, the relative homology $H_q(\mathbb{E}^{1,\leftarrow}_{p,*}(Y_k^j), \mathbb{E}^{1,\leftarrow}_{p,*}(Y_k^{j-1})) = 0$.

Proof of claim. Let δ be the unique vertex in $Y_k^j \setminus Y_k^{j-1}$. If $\delta \in \mathcal{C}_{\text{split}}(\Sigma)$ is a vertex, let $a_W(\delta)$ denote the subcomplex of $\mathcal{C}_{\text{split}}(\Sigma)$ consisting of all $\delta' \in \text{lk}(\delta)$ with $W(\delta') > W(\delta)$. Note that if $\tau \subseteq a_W(\delta)$ is a cell, the inclusion map $L_{\mathcal{A}}(\tau * \delta) \hookrightarrow L_{\mathcal{A}}(\delta)$ is an isomorphism. Hence, there is short exact sequence of chain complexes

$$0 \to a_W(\delta) \otimes H_p(L_{\mathcal{A}}(\delta); \mathbb{Z}) \to \mathbb{E}^{1,\leftarrow}_{p,*}(Y_k^{j-1}) \oplus (\delta * a_W(\delta)) \otimes H_p(L_{\mathcal{A}}(\delta); \mathbb{Z}) \to \mathbb{E}^{1,\leftarrow}_{p,*}(Y_k^j) \to 0.$$

Now, by Lemma 5.5, we have $\widetilde{H}_k(a_W(\delta)) = 0$ for $k \leq q-1$, and thus we have

- $H_k(a_W(\delta) \otimes H_p(L_{\mathcal{A}}(\delta); \mathbb{Z}) = 0$ for $1 \leq k \leq q-1$, and
- $H_k(a_W(\delta) \otimes H_p(L_{\mathcal{A}}(\delta); \mathbb{Z})) = H_p(L_{\mathcal{A}}(\delta); \mathbb{Z})$.

The complex $\delta * a_W(\delta)$ is contractible, so we have

- $H_k((\delta * a_W(\delta)) \otimes H_p(L_{\mathcal{A}}(\delta); \mathbb{Z}) = 0$ for $1 \leq k$, and
- $H_0((\delta * a_W(\delta)) \otimes H_p(L_{\mathcal{A}}(\delta); \mathbb{Z}) = 0$.

Therefore the pushforward map

$$H_k(\mathbb{E}^{1,\leftarrow}_{p,*}(Y_k^j)) \to H_k(\mathbb{E}^{1,\leftarrow}_{p,*}(Y_k^{j-1}); \mathbb{Z}))$$

is surjective for $k = q$ and bijective for $k = q-1$, so the claim holds.

Given the claim, we now have $H_q(\mathbb{E}^{1,\leftarrow}_{p,*}(X_k), \mathbb{E}^{1,\leftarrow}_{p,*}(X_{k-1})) = 0$ for $1 \leq k \leq p$. Since $H_q(\mathbb{E}^{1,\leftarrow}_{p,*}(X_0)) = 0$, we have $H_q(\mathbb{E}^{1,\leftarrow}_{p,*}(X_p)) = H_q(\mathbb{E}^{1,\leftarrow}_{p,*}) = 0$, as desired. □

We are now ready to show that $\mathfrak{a}(\mathcal{C}_{\text{split}}(\Sigma)) \geq g - 3 + \mathbb{1}_{|B_+| \geq 2} + \mathbb{1}_{|B_-| \geq 2}$.

Proof of Proposition 5.1. We will induct on the poset $\mathfrak{VerCom}$. Namely, if Σ is a vertex complement, we will assume that Proposition 5.1 holds for all $\mathcal{T} \in \mathfrak{VerCom}$ with $\mathcal{T} < \Sigma$, and show that it holds for Σ as well.

Base cases. Our base cases are given by any choice of g, $|B_+|$ and $|B_-|$ such that

$$g - 2 + \mathbb{1}_{|B_+| \geq 2} + \mathbb{1}_{|B_-| \geq 2} \geq 0.$$

This is equivalent to

$$g + \mathbb{1}_{|B_+| \geq 2} + \mathbb{1}_{|B_-| \geq 2} \geq 2,$$

and all such choices of $g, |B_+|, |B_-|$ are as follows:

- $g \geq 2$,
- $g \geq 1, |B_+| \geq 2$ or $|B_-| \geq 2$, and
- $|B_+|, |B_-| \geq 2$.

In all these cases $C_{\text{split}}(\Sigma)$ is nonempty, so the result holds.

Induction on $\mathfrak{VerCom}$. Let $\Sigma = (S, P, B_+, B_-)$ be a vertex complement such that Proposition 5.1 holds for all $\mathcal{T} < \Sigma$. Suppose without loss of generality that $|B_+| \geq |B_-|$. Choose a set Q of compact subintervals of $\partial\Sigma$ with $|Q|$ minimal such that $(S, Q) \in \mathfrak{MarkSur}$. If σ is a cell of $\mathcal{C}_{\text{split}}(\Sigma)$, let $L_{\mathcal{A}_{\text{nosep}}(\Sigma,Q)}(\sigma)$ be the full subcomplex of $\mathcal{A}_{\text{nosep}}(\Sigma, Q)$ generated by arcs disjoint from σ. Similarly, for for cells $\tau \in \mathcal{A}_{\text{nosep}}(\Sigma, Q)$, let $L_{\mathcal{C}}(\tau)$ denote the subcomplex of $\mathcal{C}_{\text{split}}(\Sigma)$ generated by cells σ with σ disjoint from τ. Let $C_{*,*}$ be the bi-cellular complex

$$C_{p,q} = \bigoplus_{\sigma \in \mathcal{A}_{\text{nosep}}(\Sigma,Q)^{(p)}} C_q(L_C(\sigma)).$$

Let $\mathbb{E}^{*;\downarrow}_{*,*}$ and $\mathbb{E}^{*;\leftarrow}_{*,*}$ be the downward and leftward bi-cellular spectral sequences respectively. It suffices to prove that the downward sequence converges to $\mathbb{Z}$ for $p + q = 0$ and converges to 0 for $0 < p + q < \dim(\mathcal{C}_{\text{split}}(\Sigma))$, and that the leftward sequence converges to $H_{p+q}(\mathcal{C}_{\text{split}}(\Sigma))$ for $0 < p + q < \dim(\mathcal{C}_{\text{split}}(\Sigma))$.

The downward spectral sequence converges to 0 *for* $0 < p + q < \dim(\mathcal{C}_{\text{split}}(\Sigma))$. On page 1, the downward spectral sequence $\mathbb{E}^{*,\downarrow}_{*,*}$ is given by

$$\mathbb{E}^{1,\downarrow}_{p,q} = \bigoplus_{\sigma \in \mathcal{A}_{\text{nosep}}(\Sigma,Q)^{(p)}} H_q(L_C(\sigma)).$$

If $|B_+| = 2$ and σ is a vertex, then $L_{\mathcal{C}}(\sigma)$ is contractible. For all other situations, the inductive hypothesis implies that the groups $\widetilde{H}_q(L_{\mathcal{C}}(\sigma))$ are trivial for σ a p-cell and $0 \leq q \leq g(\Sigma) - 3 + \mathbb{1}_{|B_+| \geq 2} + \mathbb{1}_{|B_-| \geq 2} - p$. Therefore we have

$$\mathbb{E}^{*,\downarrow}_{p,q} \Rightarrow H_{p+q}(\mathcal{A}_{\text{nosep}}(\Sigma, Q))$$

for $p+q \leq g(\Sigma)-3+\mathbb{1}_{|B_+|\geq 2}+\mathbb{1}_{|B_-|\geq 2}$. But $g-3+|B_+| \geq g(\Sigma)-3+\mathbb{1}_{|B_+|\geq 2}+\mathbb{1}_{|B_-|\geq 2}$ since we have assumed $|B_+| \geq |B_-|$. Hence Lemma 5.2, $\mathbb{E}^{*,\downarrow}_{p,q} \Rightarrow 0$ for $0 < p+q \leq g(\Sigma) - 3 + \mathbb{1}_{|B_+|\geq 2} + \mathbb{1}_{|B_-|\geq 2}$ and $\mathbb{Z}$ for $p + q = 0$.

The leftward spectral sequence converges to $H_{p+q}(\mathcal{C}_{\text{split}}(\Sigma))$. This is the content of Lemma 5.4.

Since both $\mathbb{E}^{1,\leftarrow}_{p,q}$ and $\mathbb{E}^{1,\downarrow}_{p,q}$ converge to the total homology of $C_{p,q}$, we have $\widetilde{H}_{p+q}(\mathcal{C}_{\text{split}}(\Sigma)) = 0$ for $p + q < \dim(\mathcal{C}_{\text{split}}(\Sigma))$, as desired. □

6. PL–Morse theory for cell complexes

Recall from the introduction that the proof of Theorem A proceeds by using PL–Morse theory on the complex of minimizing cycles $\mathcal{B}_{\vec{x}}(S_g)$ [BBM10]. The complex $\mathcal{B}_{\vec{x}}(S_g)$ is not a simplicial complex, and as such the results of Section 2 do not directly apply. We will address this issue here, by explaining how to use PL–Morse theory on certain types of nonsimplicial complexes. The main goal is to prove Lemma 6.6, which is a version of Lemma 2.3 that applies for CW–complexes equipped with some additional linear structure.

Locally linear cell complex. Let $\mathbb{R}^{\mathbb{N}} = \bigoplus_{n\in\mathbb{N}}\mathbb{R}$. Let $S_1, S_2 \subseteq \mathbb{R}^{\mathbb{N}}$ be two subsets. The convex hull of S_1, denoted $\mathrm{Hull}(S_1)$, is the set

$$\{t_0 s_0 + \ldots + t_n s_n : n \in \mathbb{Z}_{\geq 0}, s_0, \ldots, s_n \in S, \sum_{i=0}^{n} t_i = 1, t_i \geq 0 \text{ for all } 0 \leq i \leq n\}.$$

The convex join of S_1 and S_2, denoted $\mathrm{CJoin}(S_1, S_2)$, is the set

$$\{t_1 s_1 + t_2 s_2 : s_i \in S_i, t_1 + t_2 = 1, t_i \geq 0\}.$$

A finite dimensional cell complex X is locally linear if there is an inclusion $\iota : X \hookrightarrow \mathbb{R}^{\mathbb{N}}$ such that each cell σ of X is the convex hull of its vertices and $\partial\sigma$ is the union of the faces of $\iota(\sigma)$. We will use the notation $\mathrm{Hull}(S)$ to denote the convex hull of a set $S \subseteq \mathbb{R}^{\mathbb{N}}$. We will conflate X with its image under this map. Let $W : X^{(0)} \to \mathbb{N}$ be a PL–Morse funtion. We say that W is a *linear PL–Morse function* if W is the restriction of some linear function $\mathbb{R}^{\mathbb{N}} \to \mathbb{R}$, which by abuse of notation we will refer to as W as well. Linear PL–Morse functions have the following property.

Lemma 6.1. *Let σ be a cell of a locally linear cell complex X and let W be a linear PL–Morse function on X. The set of all points $x \in \sigma$ with $W(x)$ of maximal weight among the points of σ is a face of σ, and is the convex hull of the vertices of σ of maximal weight.*

We will say that $W : X^{(0)} \to \mathbb{N}$ is a *quasi–linear PL–Morse function* if for every cell $\sigma \subseteq X$, the convex hull of the set of all vertices $v \in \sigma^{(0)}$ with $W(v)$ maximal over all vertiex of $\sigma^{(0)}$ is a face of σ. We say that a cell $\sigma \subseteq X$ is *W–constant* if $W(v) = W(w)$ for every $v, w \in \sigma^{(0)}$. We say that W is *sharp* if the W–constant cells of X are vertices.

Descending links in locally linear cell complexes. Let X be a locally linear cell complex and let W be a quasi–linear PL–Morse function on X. Let $\sigma \subseteq X$ be a W–constant cell. We will define three different versions of the descending link $d_W(\sigma)$, and then show that each type of descending link has the same homology as each other type of descending link. Let $\mathcal{R}_\sigma$ denote the set of all cells $\rho \subseteq X$ such that

- σ is a face of ρ, and
- $W(r) \leq W(\sigma)$ for all $r \in \rho^{(0)}$, with equality if and only if $r \in \sigma$.

The variants of the descending link are as follows.

- *The facial descending link* $d_W^{\text{face}}(\sigma)$*.* Let $\rho \in \mathcal{R}_\sigma$ be a cell and let $\mathcal{T}(\rho)$ be the set of all faces $\tau \subseteq \rho$ such that $\tau \cap \sigma = \emptyset$. We define
$$d_W^{\text{face}}(\sigma) = \bigcup_{\rho \in \mathcal{R}_\sigma} \mathcal{T}(\rho).$$
- *The adjacent descending link* $d_W^{\text{adj}}(\sigma)$*.* For all $\rho \in \mathcal{R}_\sigma$, define $V(\rho)$ to be the set of all vertices $v \in \rho^{(0)}$ with $v \notin \sigma$ such that $v * w$ is an edge of X for some $w \in \sigma$. We define
$$d_W^{\text{adj}}(\sigma) = \bigcup_{\rho \in \mathcal{R}_\sigma} \text{Hull}(V(\rho)).$$
- *The total descending link* $d_W^{\text{tot}}(\sigma)$*.* This is given by
$$d_W^{\text{tot}}(\sigma) = \text{CJoin}(d_W^{\text{face}}(\sigma), d_W^{\text{adj}}(\sigma)).$$

Note that there are two inclusions
$$d_W^{\text{face}}(\sigma) \hookrightarrow d_W^{\text{tot}}(\sigma) \hookleftarrow d_W^{\text{adj}}(\sigma).$$

We have the following lemma that describes the relationships between the three types of descending links.

Lemma 6.2. *Let W be a quasi–linear PL–Morse function on a locally linear cell complex X. Let $\sigma \subseteq X$ be a W–constant cell. Then each inclusion*
$$d_W^{\text{face}}(\sigma) \hookrightarrow d_W^{\text{tot}}(\sigma) \hookleftarrow d_W^{\text{adj}}(\sigma)$$
induces an isomorphism in homology.

Proof. Let $\mathcal{R}_\sigma$ be as above. Let $\mathcal{R}_\sigma^{\max}$ denote the set of maximal cells in $\mathcal{R}_\sigma$, i.e., the set of all ρ such that there is no $\rho' \in \mathcal{R}_\sigma$ with $\rho \subsetneq \rho'$. For each choice of $\dagger = \text{face}, \text{tot}, \text{adj}$ and for each $\rho \in \mathcal{R}^{\max}$ let
$$U_\rho^\dagger = \rho \cap d_W^\dagger(\sigma).$$
Since W is quasi–linear, each set $U_\rho^\dagger$ is contractible. Furthermore, let $\rho_0, \ldots, \rho_k \in \mathcal{R}_\sigma^{\max}$. Suppose that $U_{\rho_0}^\dagger \cap \ldots \cap U_{\rho_k}^\dagger \neq \emptyset$. Then for any choice of $\dagger' = \text{tot}, \text{face}, \text{adj}$,
$$U_{\rho_0}^{\dagger'} \cap \ldots \cap U_{\rho_k}^{\dagger'} \simeq *.$$
Indeed, since $\rho_0, \ldots, \rho_k$ are cells of a locally linear complex X, their intersection $\rho = \rho_0 \cap \ldots \cap \rho_k$ must be a cell of X. Then by assumption, there must be a vertex $v \in \rho$ with $v \notin \sigma$. Hence

$V(\rho)$ as defined in the definition of d_W^{adj} must be nonempty, and thus $\mathrm{Hull}(V(\rho)) = \rho \cap d_W^{\mathrm{adj}}(\sigma)$ is contractible. Likewise, $\mathcal{T}(\rho)$ as in the definition of d_W^{face} must be nonempty, since $V(\rho) \subseteq \mathcal{T}(\rho)$, so $d_W^{\mathrm{face}}(\sigma) \cap \rho$ is nonempty. Since σ is a face of ρ because W is quasi–linear, the subcomplex $d_W^{\mathrm{face}}(\sigma) \cap \rho$ is given by taking ρ and removing a contractible subspace glued over a contractible subspace, so we conclude that $d_W^{\mathrm{face}}(\sigma) \cap \rho$ is contractible. Finally, if both of the above hold, then $\rho \cap d_W^{\mathrm{tot}}(\sigma)$ is contractible as well. Let $\mathcal{N}(U_\rho^\dagger)$ denotes the nerve of the set $\{U_\rho^\dagger\}_{\rho \in \mathcal{R}_\sigma^{\max}}$. The nerve lemma implies that we have isomorphisms in homology

$$H_*(\mathcal{N}(U_\rho^\dagger); \mathbb{Z}) \cong H_*(d_W^\dagger(\sigma); \mathbb{Z}).$$

It now remains to show that for a choice of $\dagger = \mathrm{adj}, \mathrm{face}$, the induced map of chain complexes

$$C_*(\mathcal{N}(U_\rho^\dagger); \mathbb{Z}) \to C_*(\mathcal{N}(U_\rho^{\mathrm{tot}}); \mathbb{Z})$$

induces an isomorphism in homology. In fact, we will show that the above map induces an isomorphism of chain complexes. It is clear that the induced map is an injection for either choice of $\dagger$, so it remains to prove that the natural map is a surjection.

Let $\rho_0, \ldots, \rho_n \in \mathcal{R}_\sigma^{\max}$ be a collection of cells with $U_{\rho_0}^{\mathrm{tot}} \cap \ldots \cap U_{\rho_n}^{\mathrm{tot}} \neq \emptyset$. Let $\rho = \rho_0 \cap \ldots \cap \rho_n$. ρ is a cell of X and $\sigma \neq \rho$, so there is a vertex $\tau \in X$ with $\tau \cap \sigma = \emptyset$ and $\tau \subseteq U_\rho^{\mathrm{tot}}$. Now, τ must contain a vertex adjacent to σ, since ρ is a cell. Therefore U_ρ^{adj} and U_ρ^{face} are both nonempty, so the induced map of chain complexes is surjective. □

We will say that a linear PL–Morse function W on a locally linear complex X is *n–acyclic* if $d_W^{\mathrm{tot}}(\sigma)$ is $(n - \dim(\sigma))$–acyclic for every W–constant cell $\sigma \subseteq X$. We will need one more auxiliary result.

Lemma 6.3. *Let W be a linear PL–Morse function on a locally linear complex X, let σ be a W–constant cell, and let $k = \dim(\sigma)$. If $d_W^{\mathrm{tot}}(\sigma)$ is n–acyclic for some n, then* $\mathrm{CJoin}(\partial\sigma, d_W^{\mathrm{adj}}(\sigma))$ *is $(n+k)$–acyclic.*

Proof. By Lemma 6.2, $d_W^{\mathrm{adj}}(\sigma)$ is k–acyclic. We will show that $\mathrm{CJoin}(\partial\sigma, d_W^{\mathrm{adj}})$ has the same homology groups as $\partial\sigma * d_W^{\mathrm{adj}}$. Applying Lemma 2.1 completes the proof.

Recall the set $\mathcal{R}_\sigma$ and $\mathcal{R}_\sigma^{\max}$ from Lemma 6.2. For each $\rho \in \mathcal{R}_\sigma^{\max}$, let $U_\rho = \mathrm{CJoin}(\partial\sigma, \rho \cap d_W^{\mathrm{adj}})$. Then the set

$$\mathcal{U} = \{U_\rho\}_{\rho \in \mathcal{R}_\sigma^{\max}}$$

is a cellular cover of $\mathrm{CJoin}(\partial\sigma, d_W^{\mathrm{adj}})$. Furthermore, for any $\rho_1, \ldots, \rho_n \in \mathcal{R}_\sigma^{\max}$, there is a natural map

$$\iota : \partial\sigma * (\rho_1 \cap \ldots \cap \rho_n \cap d_W^{\mathrm{adj}}) \hookrightarrow U_{\rho_1} \cap \ldots \cap U_{\rho_n}.$$

Since the set $\left\{\partial\sigma * \left(\rho \cap d_W^{\mathrm{adj}}\right)\right\}_{\rho\in\mathcal{R}_\sigma^{\max}}$ is a cover of $\partial\sigma * d_W^{\mathrm{adj}}(\sigma)$ by contractible sets, it suffices to show that ι is a homotopy equivalence for any $\rho_1, \ldots, \rho_n \in \mathcal{R}_\sigma^{\max}$.

Let $\rho_1, \ldots, \rho_n \in \mathcal{R}_\sigma^{\max}$. Let $x \in U_{\rho_1} \cap \ldots \cap U_{\rho_n}$. Since each $U_{\rho_i} \subseteq \rho_i$, we have $x \in \rho_1 \cap \ldots \cap \rho_n$. Hence $U_{\rho_1} \cap \ldots \cap U_{\rho_n}$ contains the minimal dimensional cell τ with $x \in \tau$ and $\tau \subseteq \rho_1 \cap \ldots \cap \rho_n$. Furthermore, if τ is such cell and $\tau' \subseteq \tau$ is a maximal dimensional subcell with $\tau \cap \sigma = \emptyset$, we have $\partial\sigma * \tau' \subseteq U_{\rho_1} \cap \ldots \cap U_{\rho_n}$. Then we must have $\tau' \subseteq d_W^{\mathrm{adj}}(\sigma)$ by our assumption on x. Hence $U_{\rho_1} \cap \ldots \cap U_{\rho_n} = \mathrm{CJoin}(\partial\sigma, \rho_1 \cap \ldots \cap \rho_n)$. Hence ι is a homotopy equivalence, since the source and target of ι are both either empty or contractible. Hence the natural map of chain complexes

$$\mathcal{N}(\partial\sigma * (\rho_1 \cap d_W^{\mathrm{adj}}(\sigma); \mathbb{Z}) \to \mathcal{N}(U_\rho; \mathbb{Z})$$

is an isomorphism of chain complexes. Since each intersection is contractible or empty, an application of the nerve lemma followed by Lemma 2.1 completes the proof. □

We now have the following lemma, which is similar to Lemma 2.2.

Lemma 6.4. *Let X be a finite dimensional locally linear cell complex and let W be a linear PL–Morse function. Suppose that W is n–acyclic. Then the pair $(X, M(X))$ is $(n+1)$–acyclic.*

Proof. This proceeds by the same strategy as Lemma 2.2. Let

$$X_{m,k} \coloneqq W^{-1}([0,m)) \bigcup \{\sigma \text{ a cell of } X : \max\{W|_\sigma\} \leq m, \dim(\sigma) \leq k\}.$$

If $m > 0$, then $X_{m,k}$ is built out of $X_{m-1,k}$ iteratively by attaching, for each cell $\sigma \subseteq X_{m,k}$ with $\sigma \not\subseteq X_{m-1,k}$, the following spaces in order:

- $d_W^{\mathrm{tot}}(\sigma)$ over $d_W^{\mathrm{face}}(\sigma)$, then
- $\mathrm{CJoin}(\partial\sigma, d_W^{adj}(\sigma))$ over $X_{m-1,k}$, and
- attaching a convex neighborhood of the cell σ.

The first attachment doesn't change homotopy type by Lemma 6.2. The second is given by gluing in contractible subspaces over contractible subspaces, which also induces a homotopy equivalence. The third is attaching a contractible space over an n–acyclic space by Lemma 6.3, so $H_*(X_{m,k}) \cong H_*(X_{m-1,k})$ for $* \leq n+1$.

□

We have the following lemma, which is similar to Lemma 6.4 except it deals with the case that W is sharp.

Lemma 6.5. *Let X be a finite dimensional locally linear cell complex and let W be a sharp, quasi–linear PL–Morse function. Suppose that W is n–acyclic. Then the pair $(X, M(X))$ is $(n+1)$–acyclic.*

Proof. This follows from a similar argument as Lemma 2.2. Let X_k denote the subcomplex of X generated by vertices v such that $W(v) \leq k$. Then since W is sharp, the complex X_k is constructed from X_{k-1} by attaching each $v \in X^{(0)}$ with $W(v) = k$ over the complex $d_W^{\text{face}}(v)$. By hypothesis $d_W^{\text{face}}(v)$ is n–acyclic, so the Mayer–Vietoris sequence implies that (X_k, X_{k-1}) is relatively $(n+1)$–acyclic. Since $X = \bigcup_{k\geq 0} X_k$, we have $(X, M(W))$ relatively $(n+1)$–acyclic, as desired. □

We now have the following lemma, which follows from an argument similar to the argument used in Lemma 2.3.

Lemma 6.6. *Let X be an $(n+1)$–acyclic, finite dimensional, locally linear cell complex and let W be an n–acyclic quasi–linear PL–Morse function on X. Assume that W is either sharp or linear. Then $M(W)$ is n–acyclic.*

Proof. By the long exact sequence in relative homology for the pair $(X, M(W))$ and the fact that $\widetilde{H}_k(X; \mathbb{Z}) = 0$ for $k \leq n+1$, the connecting homomorphism

$$\widetilde{H}_{k+1}(X, M(W); \mathbb{Z}) \to \widetilde{H}_k(M(W); \mathbb{Z})$$

is an isomorphism for $0 \leq k \leq n$. Hence Lemmas 6.4 and 6.5 complete the proof. □

7. The Complex of homologous curves

We now move on to the study of the complex of homologous curves $\mathcal{C}_{\vec{x}}(S_g)$. Our goal is to prove our main theorem (Theorem A), which says that $C_{\vec{x}}(S_g)$ is $(g-3)$–acyclic for $g \geq 2$. This will conclude the paper.

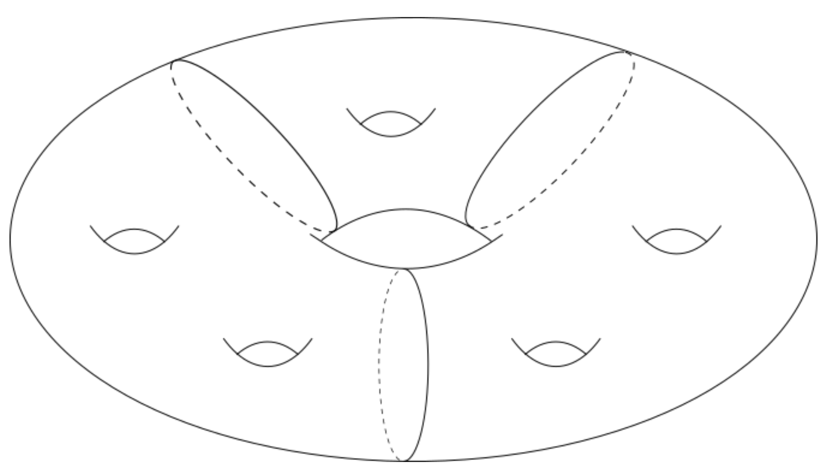

Figure 7. A 2-cell in $\mathcal{C}_{\vec{x}}(S_g)$.

The complex of homologous cycles. Let S_g be a closed, oriented surface of genus g. Let $\vec{x} \in H_1(S_g;\mathbb{Z})$ be a nonzero primitive homology class. The *complex of homologous curves* $\mathcal{C}_{\vec{x}}(S_g)$, defined by Putman [Put08], is the complex where a k–cell is an oriented multicurve $M = x_0 \sqcup \ldots \sqcup x_k$ such that $[x_i] = \vec{x}$ for every $0 \leq i \leq k$, where $[x_i]$ denotes the homology class of x_i.

The outline of the proof of Theorem A. In Section 7.1, we discuss the complex of minimizing cycles [BBM10]. We also discuss a PL-Morse function W (in the sense of Section 6) on the complex of minimizing cycles defined by Hatcher and Margalit [HM12]. The min–set of W will be $\mathcal{C}_{\vec{x}}(S_g)$. In Section 7.2, we show that an auxiliary complex called the complex of draining cycles is highly acyclic. Instances of this complex will turn out to be the descending links of the PL–Morse function W. We complete the proof of Theorem A in Section 7.3 using the PL–Morse function W and the results of Lemma 6.

7.1. **The complex of minimizing cycles.** We first define the complex of minimizing cycles, and then discuss Hatcher and Margalit's PL–Morse function W on this complex.

Let $S = S_g$ and let $\vec{x} \in H_1(S_g;\mathbb{Z})$ a primitive nonzero homology class. A *basic cycle* for $\vec{x}$ is an oriented multicurve $M = a_1 \sqcup \ldots \sqcup a_k$ such that there is a unique collection of positive integers $\lambda_1, \ldots, \lambda_k$ with

$$\vec{x} = \sum_{i=1}^{k} \lambda_i [a_i].$$

We will say that a multicurve $M = a_1 \sqcup \ldots \sqcup a_m$ is a *cycle* if

(1) each a_i is a member of a basic cycle $M' \subseteq M$, and
(2) any nontrivial linear combination of the $[a_i]$ with nonnegative $\mathbb{Z}$–coefficients is nonzero.

Remark. Condition (2) is not present in Bestvina, Bux, and Margalit's original definition. Gaifullin demonstrated that condition (2) needed to be included in the definition [Gai19]. All of Bestvina, Bux and Margalit's [BBM10] results still hold, since they implicitly assumed that condition (2) followed from (1).

Let $\mathfrak{S}$ be the set of isotopy classes of oriented nonseparating simple closed curves in S and let $\mathbb{R}^{\mathfrak{S}}$ be the real vector space consisting of finite linear combinations of elements of $\mathfrak{S}$. If M is a cycle, denote by $P_M \subseteq \mathbb{R}^{\mathfrak{S}}$ the convex hull of the basic cycles in M. Let $\mathcal{M}$ denote the set of all cycles in S_g. The complex of minimizing cycles $\mathcal{B}_{\vec{x}}(S_g)$ is the union

$$\bigcup_{M \in \mathcal{M}} P_M.$$

Observe that $\mathcal{B}_{\vec{x}}(S_g)$ is a CW-complex where the k–cells correspond to cycles $M \in \mathcal{M}$ with $\dim(P_M) = k$. Bestvina, Bux and Margalit proved that the complex $\mathcal{B}_{\vec{x}}(S)$ is contractible

[BBM10]. We will prove Theorem A by constructing a sharp, quasi-linear PL-Morse function W on $\mathcal{B}_{\vec{x}}(S_g)$ that is $(g-3)$–acyclic and has min-set equal to $\mathcal{C}_{\vec{x}}(S)$.

The PL-Morse function. We will use the PL-Morse function on $\mathcal{B}_{\vec{x}}(S_g)$ originally defined by Hatcher and Margalit [HM12]. Let v be a vertex in $\mathcal{B}_{\vec{x}}(S)$ represented by a basic cycle $M = a_0 \sqcup \ldots \sqcup a_m$. By definition, $\vec{x} = \sum_{i \leq m} \lambda_i [a_i]$ for some unique positive integers λ_i. We define

$$W(v) = \sum_{i \leq m} \lambda_i.$$

The min-set $M(W)$. There are no vertices $v \in \mathcal{B}_{\vec{x}}(S_g)$ with $W(v) = 0$. In fact, the lowest weight vertices are given by all v with $W(v) = 1$. The subcomplex of $\mathcal{B}_{\vec{x}}(S)$ generated by v with $W(v) = 1$ is precisely the complex of homologous curves $\mathcal{C}_{\vec{x}}(S)$, so by we will let $M(W)$ denote the subcomplex of $\mathcal{B}_{\vec{x}}(S)$ generated by vertices v with $W(v) = 1$.

7.2. **The complex of draining cycles.** Our goal in this section is to prove Lemma 7.1, which says that a certain complex called the complex of draining cycles is highly acyclic. This complex is a generalization of the complex of splitting curves introduced in Section 5. We will leverage Lemma 7.1 to prove that the PL-Morse function W is $(g-3)$–acyclic.

The poset of partial cobordisms. We begin by defining a certain poset of labeled surfaces. Let $S = S_g^b$. We say that S is a *partial cobordism* if it comes equipped with an orientation of some of the boundary components of S. We let B_+ be the set of boundary components oriented inward (according to the right hand rule) and B_- the set oriented outward. We will refer to such an S as a *partial cobordism from B_+ to B_-*. We will denote a partial cobordism by $\Sigma = (S, B_+, B_-, B_0)$, where $B_0 = \pi_0(\partial S) \setminus (B_+ \cup B_-)$. We will refer to S as the underlying surface of Σ. We will say that a partial cobordism is

- *draining* if $|B_+| > |B_-|$,
- *balanced* if $|B_+| = |B_-|$, and
- *flooding* if $|B_+| < |B_-|$.

The set of partial cobordisms forms a poset, denoted $\mathfrak{ParCob}$. The order is given by $\Sigma < \Sigma'$ if either $g(\Sigma) < g(\Sigma')$, or $g(\Sigma) = g(\Sigma')$ and $|B_+| < |B'_+|$.

Cutting multicurves in partial cobordisms. Let $\Sigma = (S, B_+, B_-, B_0)$ be a partial cobordism. Let $M \subseteq S$ be an oriented multicurve such that $|\pi_0(S \searrow M)| = 2$. Then each connected component of $S \searrow M$ naturally has the structure of a partial cobordism, since the boundary components of $S \searrow M$ corresponding to curves in M come equipped with orientations.

The complex of draining cycles. Let $S = S_g^b$ and let $\Sigma = (S, B_+, B_-, B_0)$ be a partial cobordism. The complex of draining cycles, denoted $\mathcal{C}_{\mathrm{dr}}(\Sigma)$, is a complex associated to any $\Sigma \in \mathfrak{ParCob}$. A vertex of $\mathcal{C}_{\mathrm{dr}}(\Sigma)$ is an oriented multicurve $M \subseteq S$ such that

- $|\pi_0(S \searrow M)| = 2$,
- one connected component of $S \searrow M$ is a partial cobordism Σ_M from a subset of B_+ to a union of M and a subset of B_-,
- the partial cobordism Σ_M is draining, and
- any nontrivial linear combination of the homology classes represented by curves in $M \cup B_+$ with nonnegative integer coefficients is nonzero.

Note that the partial cobordism Σ_M is unique. Two examples of vertices in this complex can be found in Figure 8. We now define the higher-dimensional cells in the complex $\mathcal{C}_{\mathrm{dr}}(\Sigma)$. Let $\mathfrak{S}_\Sigma$

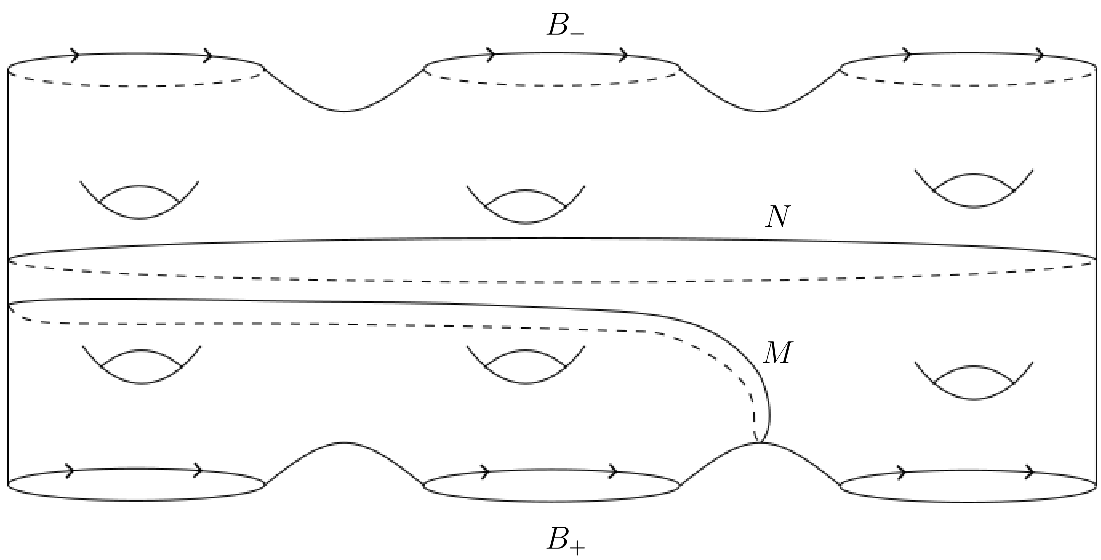

FIGURE 8. Two vertices M and N in the complex of draining cycles.

be set of oriented isotopy classes of essential simple closed curves on Σ. We say that an oriented multicurve $M \subseteq S$ is *representative* if

- each curve in M is contained in a vertex of $\mathcal{C}_{\mathrm{dr}}(\Sigma)$ supported on M,
- if $N \subseteq M$ is a multicurve such that $|\pi_0(S \searrow N)| = 2$ and at least one connected component Σ_N of $S \searrow N$ is a partial cobordism from a subset of B_+ to a union of N and a subset of B_-, then Σ_N is draining, and

- any nontrivial linear combination of the homology classes represented by curves in $M \cup B_+$ with nonnegative integer coefficients is nonzero in $H_1(S;\mathbb{Z})$, where S is the underlying surface of Σ.

Let $\mathfrak{S}_\Sigma$ denote the set of all oriented simple closed curves in S. Let $\mathcal{M}$ denote the set of representative multicurves in Σ. Each vertex $M \in \mathcal{M}$ corresponds to the point $a_1 + \ldots + a_k \in \mathbb{R}^{\mathcal{S}}$ with $M = a_1 \sqcup \ldots \sqcup a_k$. We will henceforth identify a vertex $M \in \mathcal{M}$ with this point in $\mathbb{R}^{\mathcal{S}}$. For an arbitrary $M \in \mathcal{M}$, let $P_M \subseteq \mathbb{R}^{\mathcal{S}_\Sigma}$ be the convex hull of the vertices supported in M. The complex $\mathcal{C}_{\mathrm{dr}}(\Sigma)$ is the union

$$\bigcup_{M \in \mathcal{M}} P_M.$$

Observe that, by construction, $\mathcal{C}_{\mathrm{dr}}(\Sigma)$ is a locally linear complex. The main goal of Section 7.2 is to prove the following lemma.

Lemma 7.1. *Let $\Sigma = (S_g^b, B_+, B_-, B_0)$ be a partial cobordism with $|B_+|, |B_-| \geq 2$ and $|B_+| \leq |B_-|$. The complex $\mathcal{C}_{\mathrm{dr}}(\Sigma)$ is $(g-3+|B_+|)$–acyclic.*

In order to prove Lemma 7.1, we will need the following auxiliary result.

Lemma 7.2. *Let $\Sigma = (S_g^b, B_+, B_-, B_0)$ be a partial cobordism with $|B_+|, |B_-| \geq 2$, $|B_0| \geq 1$. Let $p \in B_0$ be a boundary component and let Σ' be the partial cobordism given by gluing a disk along p. The induced map $\mathcal{C}_{\mathrm{dr}}(\Sigma) \to \mathcal{C}_{\mathrm{dr}}(\Sigma')$ is a homotopy equivalence.*

Proof. Let δ be an arc in S with one endpoint in p and another endpoint in $p' \in B_+$. Assume that δ is oriented from p to p'. Let Surg_δ be defined as in Section 4. Surg_δ extends to a function on multicurves M by surgering whichever $c \subseteq M$ has a point of intersection closest along α to p', and discarding any extra copies of isotopic curves. Now, if $M \in \mathcal{C}_{\mathrm{dr}}(\Sigma)$ is a vertex, $\mathrm{Surg}_\delta(M)$ is also a vertex of $\mathcal{C}_{\mathrm{dr}}(\Sigma)$. Furthermore, the subcomplex of $\mathcal{C}_{\mathrm{dr}}(\Sigma)$ consisting of multicurves M disjoint from β is isomorphic to $\mathcal{C}_{\mathrm{dr}}(\Sigma)$, since $\Sigma \setminus \beta$ is naturally identified as partial cobordism with Σ'. Then by a similar argument to the proof of Proposition 4.1, there is a homotopy equivalence $\mathcal{C}_{\mathrm{dr}}(\Sigma) \to \mathcal{C}_{\mathrm{dr}}(\Sigma')$. Since Surg_β fixes any M disjoint from β, the proof is complete. ☐

We now show that $\mathcal{C}_{\mathrm{dr}}(\Sigma)$ is $(g-3+|B_+|)$–acyclic.

Proof of Lemma 7.1. We will induct on the poset of partial cobordism $\mathfrak{ParCob}$.

Base cases. The base cases are given by partial cobordisms with $|B_+| = 2$ and $|B_0| = 0$. In this case, $\mathcal{C}_{\mathrm{dr}}(\Sigma)$ is identified with the complex of splitting curves. Lemma 5.1 says that when $|B_+| = 2$ and $|B_-| \geq 2$, the complex of splitting curves is $(g-1)$–acyclic, so the lemma holds.

Induction on $\mathfrak{ParCob}$. Let $\Sigma = (S_g^b, B_+, B_-, B_0) \in \mathfrak{ParCob}$ with $|B_+| \geq 3$, $|B_-| \geq 2$ and $|B_+| \leq |B_-|$ such that the lemma holds for all $\mathcal{T} < \Sigma$ satisfying the hypotheses of the lemma.

Lemma 7.2 says that we can fill in boundary components of B_0 with disks without changing the homotopy type of $\mathcal{C}_{\text{dr}}(\Sigma)$, so we may assume that $|B_0| = 0$. Let $\Sigma' = (S_g^b, B'_+, B_-, B_0)$ be the partial cobordism given by forgetting the orientation of a component $b \in B_+$, so $b \in B_0$. Let $\Sigma'' = (S_g^{b-1}, B'_+, B_-)$ be the partial cobordism given by filling in b with a disk. By Lemma 7.2, there is a homotopy equivalence

$$\mathcal{C}_{\text{dr}}(\Sigma') \simeq \mathcal{C}_{\text{dr}}(\Sigma'').$$

By the inductive hypothesis, $\mathcal{C}_{\text{dr}}(\Sigma'')$ is $(g - 4 + |B_+|)$–acyclic. Hence it suffices to show that $\mathfrak{a}(\mathcal{C}_{\text{dr}}(\Sigma')) + 1 \leq \mathfrak{a}(\mathcal{C}_{\text{dr}}(\Sigma))$. We do this in three steps:

(1) We add to $\mathcal{C}_{\text{dr}}(\Sigma')$ any vertices $M \in \mathcal{C}_{\text{dr}}(\Sigma)$ that become inessential when b is capped with a disk, and show that this strictly increases acyclicity by 1.
(2) We add vertices $M \in \mathcal{C}_{\text{dr}}(\Sigma)$ such that M is a single simple closed curve, and show that this does not decrease acyclicity.
(3) We add in the remaining vertices and show again that acyclicity does not decrease.

Adding curves that become inessential when b is filled in with a disk. Let K be the set of vertices in $\mathcal{C}_{\text{dr}}(\Sigma)$ such that at least one curve becomes inessential when b is filled in with a disk. The set K is a discrete set in the sense that no two vertices of K share an edge, and every element of K is a single curve γ such that γ surrounds b and one other boundary component in B_+. Let $\mathcal{C}_{\text{dr}}(\Sigma', K)$ denote the full subcomplex of $\mathcal{C}_{\text{dr}}(\Sigma)$ generated by $\mathcal{C}_{\text{dr}}(\Sigma')$ and K. An example of a curve in K can be found in Figure 9.
By Lemma 7.2 and Lemma 4.3, there is a homotopy equivalence

$$\mathcal{C}_{\text{dr}}(\Sigma', K) \simeq \mathcal{C}_{\text{dr}}(\Sigma') * K.$$

Hence $\mathfrak{a}(\mathcal{C}_{\text{dr}}(\Sigma', K)) \geq (g - 3 + |B_+|)$ by Lemma 2.1 and the inductive hypothesis.

Adding in single curves. Let $C^1_{\text{dr}}(\Sigma', K)$ denote the subcomplex of $\mathcal{C}_{\text{dr}}(\Sigma)$ consisting of multicurves M such that every vertex $N \subseteq M$ either has $N \in \mathcal{C}_{\text{dr}}(\Sigma', K)$ or N a single curve. For any vertex $M \in C^1_{\text{dr}}(\Sigma', K)$ with $M \notin C_{\text{dr}}(\Sigma', K)$, let Σ_M denote the associated cobordism from the definition of the complex of draining cycles. Let $\mathcal{T}_M$ denote the other connected component of $S_g^b \setminus M$, which is naturally a cobordism. An example of M, Σ_M and $\mathcal{T}_M$ can be found in Figure 10.

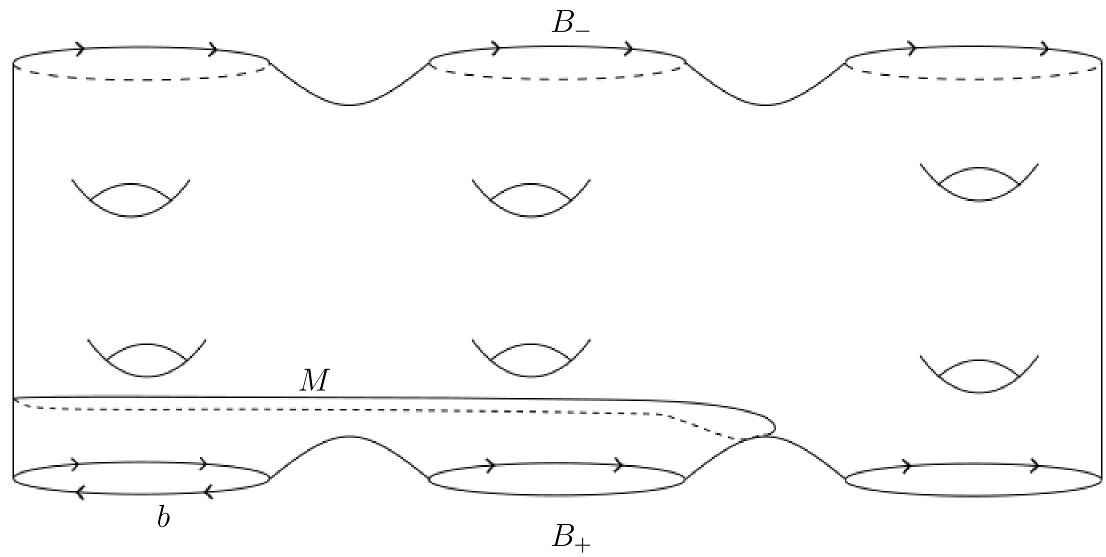

FIGURE 9. M is a vertex in in K.

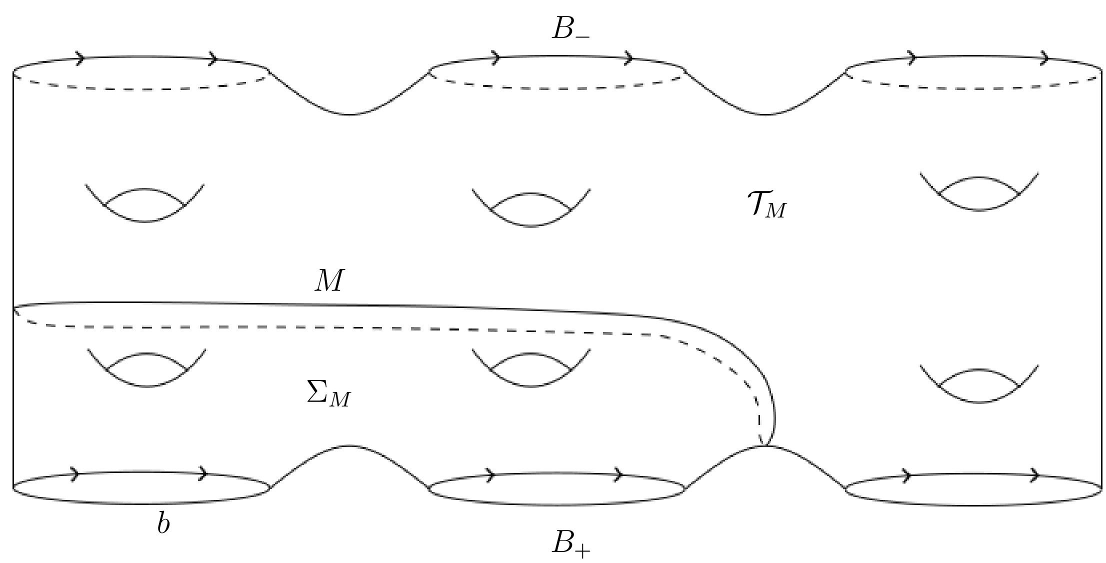

FIGURE 10. A curve $M \in C^1_{\mathrm{dr}}(\Sigma', K)$. The top boundary components are in B_-, the bottom components are in B_+.

For M a vertex of $\mathcal{C}^1_{\mathrm{dr}}(\Sigma', K)$, let

$$W(M) = \begin{cases} 0 & \text{if } M \in \mathcal{C}_{\mathrm{dr}}(\Sigma', K) \\ g(\Sigma_M) & \text{otherwise.} \end{cases}$$

We will show that W is a $(g-4+|B_+|)$–acyclic PL–Morse function. Since $\mathcal{C}_{\mathrm{dr}}(\Sigma', K)$ is at least $(g-3+|B_+|)$–acyclic and W is sharp (as in Section 6), Lemma 6.6 will imply that $\mathcal{C}^1_{\mathrm{dr}}(\Sigma', K)$ is $(g-3+|B_+|)$–acyclic.

Let $M \in \mathcal{C}^1_{\mathrm{dr}}(\Sigma', K)$ be a positive weight vertex and let $N \in d^{\mathrm{adj}}_W(M)$ be a vertex. Observe that N is supported on either $\mathcal{T}_M$ or Σ_M. If N is supported on $\mathcal{T}_M$, then $N \in \mathcal{C}_{\mathrm{dr}}(\mathcal{T}_M)$. Otherwise, $N \in \mathcal{C}_{\mathrm{dr}}(\Sigma_M)$, which implies that N is a single curve. Hence there is a canonical isomorphism

$$d_W(M) = \mathcal{C}_{\mathrm{split}}(\Sigma_M) * \mathcal{C}_{\mathrm{dr}}(\mathcal{T}_M).$$

By Proposition 5.1, $\mathfrak{a}(\mathcal{C}_{\mathrm{split}}(\Sigma_M)) = g(\Sigma_M) - 2$. Then by the inductive hypothesis, we have

$$\mathfrak{a}(\mathcal{C}_{\mathrm{dr}}(\mathcal{T}_M)) = g(\mathcal{T}_M) - 3 + |B_+| - 1 = g(\mathcal{T}_M) - 4 + |B_+|.$$

Hence by Lemma 2.1, we have

$$\mathfrak{a}(d_W(M)) = g(\Sigma_M) - 2 + g(\mathcal{T}_M) - 4 + |B_+| + 2 = g(\Sigma) - 4 + |B_+|$$

so this step is complete.

Adding other vertices. We will now add in vertices in $\mathcal{C}_{\mathrm{dr}}(\Sigma) \setminus \mathcal{C}^1_{\mathrm{dr}}(\Sigma', K)$. Let M be a vertex of $\mathcal{C}_{\mathrm{dr}}(\Sigma)$ and let Σ_M be connected component of $\Sigma \searrow M$ that realizes M as a vertex of $\mathcal{C}_{\mathrm{dr}}(\Sigma)$. Let

$$W(M) = \begin{cases} 0 & \text{if } M \in \mathcal{C}^1_{\mathrm{dr}}(\Sigma', K) \\ |\chi(\Sigma_M)| & \text{otherwise.} \end{cases}$$

We will show that W is a $(g-4-|B_+|)$–acyclic PL–Morse function. Since W is sharp, an application of Lemma 6.6 will complete the proof. Let M be a vertex in $\mathcal{C}_{\mathrm{dr}}(\Sigma)$ of positive weight. Observe that the cobordism $\Sigma_M = (S_M, B^M_+, B^M_-)$ satisfies the following two properties:

- $|B^M_+| = |B^M_-| + 1$ and
- $p \in B^M_+$

since otherwise M would be a vertex of $\mathcal{C}_{\mathrm{dr}}(\Sigma', K)$. Let $\mathcal{T}_M$ be the connected component of $\Sigma \searrow M$ which is not Σ_M. An example of such a vertex M can be found in Figure 11.

The surface $\mathcal{T}_M$ is naturally a cobordism from a union of M and a subset of B_+ to a subset of B_-. We will show that $d^{\mathrm{adj}}_W(M) \cong \mathcal{C}_{\mathrm{dr}}(\Sigma_M) * \mathcal{C}_{\mathrm{dr}}(\mathcal{T}_M)$. Let $N \in d^{\mathrm{adj}}_W(M)$ be a vertex. Let $P = M \cup N$. By definition N and M are adjacent, so N is either supported on Σ_M or $\mathcal{T}_M$. If $N \subseteq \Sigma_M$, then $N \in \mathcal{C}_{\mathrm{dr}}(\Sigma_M)$. Any vertex $P \in \mathcal{C}_{\mathrm{dr}}(\Sigma_M)$ must have $W(P) < W(M)$, so $d^{\mathrm{adj}}_W(M) \cap \mathcal{C}_{\mathrm{dr}}(\Sigma_M) = \mathcal{C}_{\mathrm{dr}}(\Sigma_M)$. Otherwise, if $N \subseteq \mathcal{T}_M$, then $W(N) = 0$. Therefore the partial cobordism Σ_N realizing N as draining is still draining even if p is filled in with a disk. Then since $\Sigma_N \supseteq \Sigma_M$, we have $N \in \mathcal{C}_{\mathrm{dr}}(\mathcal{T}_M)$.

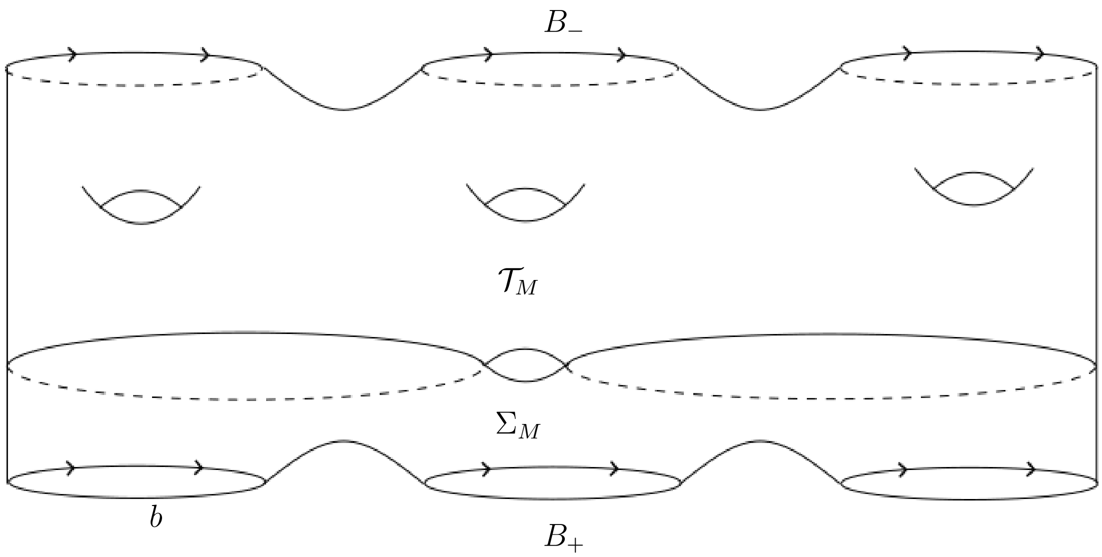

FIGURE 11. A vertex $M \in C^1_{\mathrm{dr}}(\Sigma', K)$. The top boundary components are in B_-, the bottom components are in B_+.

Hence $d_W(M)$ is at least $(g(\Sigma_M) - 3 + |B_+(\Sigma_M)| + g(\mathcal{T}_M) - 3 + |M| + |\mathcal{T}_M \cap B_+| + 2)$–acyclic by the inductive hypothesis and Lemma 2.1. But then $g(\mathcal{T}_M) + g(\Sigma_M) + |M| = g$ and $|B_+^M| + |\mathcal{T}_M \cap B_+| = |B_+|$, so $d_W(M)$ is at least $(g-4+|B_+|)$–acyclic. Therefore since $\mathcal{C}^1_{\mathrm{dr}}(\Sigma', K)$ is $(g-3+|B_+|)$–acyclic, Lemma 6.6 implies that $\mathcal{C}_{\mathrm{dr}}(\Sigma)$ is $(g-3+|B_+|)$–acyclic, which completes the proof. □

7.3. **Completing the proof.** We are now almost ready to prove Theorem A, which will conclude the paper. We will prove in Lemma 7.3 that the PL–Morse function W defined in Section 7.1 is a linear PL–Morse function. We then prove in Lemma 7.5 that the descending links of W–constant k–cells are given by joins of the complex of draining cycles. We conclude by proving Theorem A.

Lemma 7.3. *The PL–Morse function W is a linear PL–Morse function.*

Proof. Let $\sigma \subseteq \mathcal{B}_{\vec{x}}(S_g)$ be a cell corresponding to a multicurve M. By definition, σ is the convex hull of its vertices $v_0, \ldots, v_m$. Hence every point $v \in \sigma$ is a linear combination

$$v = \sum_{i=0}^{m} t_i v_i$$

with $\sum_{i=0}^{m} t_i = 1$. Set

$$W^\sigma(v) = \sum_{i=0}^{m} t_i W(v_i).$$

It suffices to show that W is well-defined on points in σ, i.e., independent of the choice of coefficients t_i. To see this, suppose that

$$v = \sum_{i=0}^{m} t_i' v_i$$

is another linear combination with $\sum_{i=0}^{m} t_i' = 1$. If $a_0, \ldots, a_n$ are the underlying simple closed curves in the multicurve M corresponding to σ, then each v_i is by definition a formal sum

$$v_i = \sum_{j=0}^{n} \lambda_{i,j} a_j$$

such that

$$\vec{x} = \sum_{j=0}^{n} \lambda_{i,j} [a_j]$$

in $H_1(S_g; \mathbb{Z})$. Since

$$\sum_{i=0}^{m} t_i v_i = \sum_{i=0}^{m} t_i' v_i$$

we have a relation

$$\sum_{i=0}^{m} t_i \sum_{j=0}^{n} \lambda_{i,j} = \sum_{i=0}^{m} t_i' \sum_{j=0}^{n} \lambda_{i,j}.$$

But $W(v_i) = \sum_{j=0}^{n} \lambda_{i,j}$, so this relation is

$$\sum_{i=0}^{m} t_i W(v_i) = \sum_{i=0}^{m} t_i' W(v_i)$$

so the claim holds. □

We now describe the W–constant cells of $\mathcal{B}_{\vec{x}}(S_g)$. If $\sigma \subseteq \mathcal{B}_{\vec{x}}(S_g)$ is a cell, we denote by $\sigma_{\max}$ the convex hull of the vertices in σ of maximal weight.

Lemma 7.4. *Let σ be a W–constant cell of $\mathcal{B}_{\vec{x}}(S_g)$. Let M be the oriented multicurve corresponding to σ and let $\Sigma_0 \sqcup \ldots \sqcup \Sigma_k$ be the connected components of $S_g \setminus M$ equipped with the appropriate partial cobordism structure. Each partial cobordism Σ_i is balanced.*

Proof. Suppose that some cobordism Σ_i is not balanced. Σ_i is a cobordism between vertices v and w of σ. If Σ_i is unbalanced, then $W(v) \neq W(w)$ which contradicts the assumption that σ is W–constant. □

We now explicitly compute $d_W(M)$ in the case that M is a W–constant k–cell.

Lemma 7.5. *Let $M \subseteq \mathcal{B}_{\vec{x}}(S_g)$ be a W–constant k–cell. Let $\Sigma_0, \ldots, \Sigma_k$ be the connected components of $S_g \setminus M$. Then the natural inclusion*

$$C_{\mathrm{dr}}(\Sigma_0) * \ldots * C_{\mathrm{dr}}(\Sigma_k) \to d_W^{\mathrm{adj}}(M)$$

is an isomorphism.

Proof. Let N be a vertex in $d_W^{\mathrm{adj}}(M)$. Then $N \searrow M$ is contained in some connected component Σ of $S_g \searrow M$, which is naturally a balanced cobordism. Let $\mathcal{T} \subseteq \Sigma$ be the cobordism between a subset of $B_+(\Sigma)$ and a union of $N \setminus M$ and a subset of $B_-(\Sigma)$. Let $P \subseteq M$ be a vertex containing $B_+(\Sigma)$, so N is adjacent to P. Then the cobordism $\mathcal{T}$ must be draining. Indeed, suppose that

$$\sum_{p \in P} \lambda_p [p] = \vec{x}$$

is the linear relation corresponding to P. Suppose that $\mathcal{T}$ realizes a relation in $H_1(S_g; \mathbb{Z})$ of the form

$$\sum_{p \in P'} [p] = \sum_{n \in N'} [n]$$

for subsets $P' \subseteq P$, $N' \subseteq N$. Then there is a linear relation in $H_1(S_g; \mathbb{Z})$ supported on N given by

$$\sum_{p \in P} \lambda_p [P] - \min_{p \in P'} \lambda_p \left(\sum_{p \in P'} [p] \right) + \min_{p \in P'} \lambda_p \left(\sum_{n \in N'} [n] \right) = \vec{x}$$

Note that this relation is a nonnegative linear combination of curves in N, since this relation is another nonnegative integral relation on curves contained in $N \cup P$ and N and P are the endpoints of an edge. Then since $W(N) < W(P)$, we have

$$W(N) = W(P) - \min_{p \in P'} \lambda_p |P'| + \min_{p \in P'} \lambda_p |N'| = W(P) - \min_{p \in P'} (|P| - |N'|).$$

Hence $|N'| < |P'|$, so $\mathcal{T}$ is draining as desired. □

We now conclude Section 7 by proving that $\mathcal{C}_{\vec{x}}(S_g)$ is $(g-3)$–acyclic.

Proof of Theorem A. Let M be a multicurve that represents a W–constant k-cell σ. Label the connected components of $S \searrow M$ by $\Sigma_0 \sqcup \ldots \sqcup \Sigma_k$. By Lemma 7.5, we have

$$d_W^{\mathrm{adj}}(M) = \mathcal{C}_{\mathrm{dr}}(\Sigma_0) * \ldots * \mathcal{C}_{\mathrm{dr}}(\Sigma_k),$$

where $\Sigma_0, \ldots, \Sigma_k$ are the connected components of $S_g \searrow M$. Hence $d_W(\sigma)$ is at least $(g-3-k)$–acyclic by Lemma 7.1 and Lemma 2.1. Now, W is a linear PL–Morse function by Lemma 7.3. Hence W is a $(g-3)$–acyclic linear PL-Morse function, so Theorem A follows by Lemma 6.6 and Bestvina, Bux and Margalit's theorem that $\mathcal{B}_{\vec{x}}(S_g)$ is contractible [BBM10, Theorem E]. □

References

[BBM10] Mladen Bestvina, Kai-Uwe Bux, and Dan Margalit. The dimension of the Torelli group. *Journal of the American Mathematical Society*, 23(1), 2010.

[BBP23] Tara Brendle, Nathan Broaddus, and Andrew Putman. The high-dimensional cohomology of the moduli space of curves with level structures II: punctures and boundary. *Isreali J. Math (To appear)*, 2023.

[Bes] Mladen Bestvina. PL morse theory minicourse. `https://www.math.utah.edu/~bestvina/eprints/minicourse.pdf`. Accessed: 2020-02-26.

[BM06] Robert W. Bell and Dan Margalit. Braid groups and the co-Hopfian property. *J. Algebra*, 303(1):275–294, 2006.

[Bro94] Kenneth S. Brown. *Cohomology of groups*. Graduate Texts in Mathematics, 87. Springer-Verlag, New York, 1994.

[Dug] Daniel Dugger. A primer on homotopy colimits. https://pages.uoregon.edu/ddugger/hocolim.pdf.

[FM12] Benson Farb and Dan Margalit. *A Primer on Mapping Class Groups*. Princeton University Press, Princeton, 2012.

[Gai19] Alexander A. Gaifullin. On infinitely generated homology of Torelli groups, 2019.

[GGV18] Jonah Gaster, Joshua Evan Greene, and Nicholas G. Vlamis. Coloring curves on surfaces. *Forum Math. Sigma*, 6:Paper No. e17, 42, 2018.

[Hat91] Allen Hatcher. On triangulations of surfaces. *Topology Appl.*, 40(2):189–194, 1991.

[HM12] Allen Hatcher and Dan Margalit. Generating the Torelli group. *Enseign. Math. (2)*, 58(1-2):165–188, 2012.

[Joh83] Dennis Johnson. The structure of the Torelli group. I. A finite set of generators for T. *Ann. of Math. (2)*, 118(3):423–442, 1983.

[MvdK02] B. Mirzaii and W. van der Kallen. Homology stability for unitary groups. *Doc. Math.*, 7:143–166, 2002.

[Put07] Andrew Putman. Cutting and pasting in the Torelli group. *Geometry & Topology*, 11(2):829–865, 2007.

[Put08] Andrew Putman. A note on the connectivity of certain complexes associated to surfaces. *Enseign. Math. (2)*, 54(3-4):287–301, 2008.

[Put12] Andrew Putman. Small generating sets for the Torelli group. *Geom. Topol.*, 16(1):111–125, 2012.

[Sch15] Saul Schleimer. `http://homepages.warwick.ac.uk/~masgar/Maths/notes.pdf`, 2015.

Skiles Classroom Building, Georgia Institute of Technology, Atlanta, GA 30332
Email address: `dminahan6@gatech.edu`