The complex of cycles on a surface (after Bestvina–Bux–Margalit)

Andrew Putman

Abstract

The complex of cycles on a surface is a cell complex that encodes all the ways that an element of first homology can be written as an embedded cycle. It was introduced by Bestvina–Bux–Margalit and plays an important role in their calculation of the cohomological dimension of the Torelli group. We give a detailed proof that this complex is contractible, expanding upon one of the proofs given by Bestvina–Bux–Margalit.

Contents

1 Introduction

This note is devoted to the *complex of reduced cycles* on a surface, which is a space that encodes all the ways that a fixed homology class can be written as a cycle. This complex was introduced by Bestvina–Bux–Margalit [1], who used it to calculate the cohomological dimension of the Torelli group and to give a topological proof of a theorem of Mess [5] that says that the genus 2 Torelli group is an infinitely generated free group. The main result in this note is Theorem 5.1, which asserts that the complex of reduced cycles is contractible. Our proof follows the "second proof" of this fact from [1]. See [4] for an alternate exposition of it.

The outline of this note is as follows. In §2, we define the complex of cycles. This complex is made up of certain "cells", and in §3 we discuss some basic properties of these cells. Next, in §4 we discuss some preliminary results needed to prove contractibility. Finally, in §5 we prove that the complex is contractible.

Throughout this note, Σ is a closed surface and $x \in \mathrm{H}_1(\Sigma;\mathbb{Z})$ is a fixed primitive element.

2 Basic definitions

We will first define the complex of cycles as a set and then discuss its topology.

Multicurves and weighted multicurves. An *oriented multicurve* γ on Σ is an unordered collection $\gamma_1 \cup \cdots \cup \gamma_k$ of disjoint oriented nonnullhomotopic simple closed curves which are pairwise non-homotopic (as unoriented curves, i.e. we do not allow one of the γ_i to be homotopic to another γ_j but with a reversed orientation). We will not distinguish between homotopic oriented multicurves. A *submulticurve* of an oriented multicurve γ is an oriented multicurve each of whose curves is also a curve in γ.

A *weighted oriented multicurve* is a formal expression $c_1\gamma_1 + \cdots + c_k\gamma_k$ with $\gamma := \gamma_1 \cup \cdots \cup \gamma_k$ an oriented multicurve and $c_1, \ldots, c_k \in \mathbb{R}$. The ordering of the γ_i in this expression does not matter. The number c_i is the *weight* of γ_i and the *support* of $c_1\gamma_1 + \cdots + c_k\gamma_k$ is the submulticurve of γ composed of all of the γ_i whose weights are nonzero. We will identify two weighted oriented multicurves that differ by inserting or deleting oriented curves of weight 0. The homology class *represented* by $c_1\gamma_1 + \cdots + c_k\gamma_k$ is $c_1[\gamma_1] + \cdots + c_k[\gamma_k] \in \mathrm{H}_1(\Sigma;\mathbb{Z})$. We will call $c_1\gamma_1 + \cdots + c_k\gamma_k$ a *positively weighted oriented multicurve* if all the weights c_i are nonnegative.

Complex of unreduced cycles as a set. The *complex of unreduced cycles*, denoted $\widehat{\mathcal{C}}_x(\Sigma)$, is the set of positively weighted oriented multicurves representing the fixed primitive homology class x. We will soon define a topology on $\widehat{\mathcal{C}}_x(\Sigma)$. Intuitively, one moves around in this topology by continuously varying the weights in positively weighted oriented multicurves while keeping the represented homology class constant. When one of the weights goes to 0, that curve disappears.

Cells. Let γ be some oriented multicurve on Σ. The *cell* associated to γ, denoted $\mathcal{X}_x(\gamma)$, is the subset of $\widehat{\mathcal{C}}_x(\Sigma)$ consisting of positively weighted oriented multicurves representing x whose support is a submulticurve of γ. We will say that $\mathcal{X}_x(\gamma)$ is *nondegenerate* if it contains a positively weighted

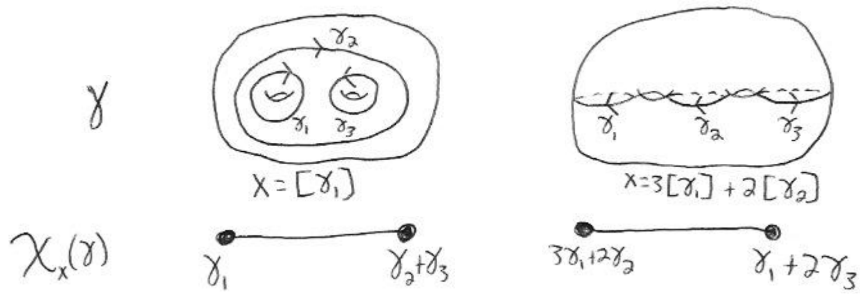

Figure 1: Example of multicurves γ whose associated cells $\mathcal{X}_x(\gamma)$ are compact 1-dimensional polyhedra. Under each multicurve is how to write x as a linear combination of the multicurves. The points in the interior of the left edge are $t\gamma_1 + s\gamma_2 + s\gamma_2$ with $s, t \geq 0$ and $t + s = 1$. The points in the interior of the right edge are $(3t + s)\gamma_1 + 2t\gamma_2 + (2t + 2s)\gamma_3$ with $s, t \geq 0$ and $s + t = 1$.

oriented multicurve whose support is equal to γ. Writing $\gamma = \gamma_1 \cup \cdots \cup \gamma_k$, there is an inclusion of $\mathcal{X}_x(\gamma)$ into $\mathbb{R}^k_{\geq 0}$ that takes $c_1\gamma_1 + \cdots + c_k\gamma_k$ to $(c_1, \ldots, c_k)$. This inclusion defines a topology on $\mathcal{X}_x(\gamma)$; in fact, it endows $\mathcal{X}_x(\gamma)$ with the structure of a (not necessarily compact) polyhedron, possibly empty. This structure does not depend the ordering of the γ_i. If γ' is a submulticurve of γ, then $\mathcal{X}_x(\gamma')$ is in a natural way a subpolyhedron of $\mathcal{X}_x(\gamma)$. See Figures 1–3 for some examples of cells.

Topology on complex of unreduced cycles. If γ is an oriented multicurve on Σ, then $\mathcal{X}_x(\gamma)$ can be regarded as a subset of $\widehat{\mathcal{C}}_x(\Sigma)$. We will give $\widehat{\mathcal{C}}_x(\Sigma)$ the weak topology with regards to the $\mathcal{X}_x(\gamma)$. In other words, a set $U \subset \widehat{\mathcal{C}}_x(\Sigma)$ is open if and only if $U \cap \mathcal{X}_x(\gamma)$ is open for all oriented multicurves γ.

Complex of reduced cycles. We will say that a cell $\mathcal{X}_x(\gamma)$ is *reduced* if it is compact. Below in Lemma 3.4 we will give an easy-to-check characterization of when a cell is reduced. We will call a positively weighted oriented multicurve $c \in \widehat{\mathcal{C}}_x(\Sigma)$ with support γ *reduced* if $\mathcal{X}_x(\gamma)$ is reduced. The *complex of reduced cycles*, denoted $\mathcal{C}_x(\Sigma)$, is the subset of $\widehat{\mathcal{C}}_x(\Sigma)$ consisting of reduced positively weighted oriented multicurves representing x. The reduced cells endow $\mathcal{C}_x(\Sigma)$ with the structure of a polyhedral complex.

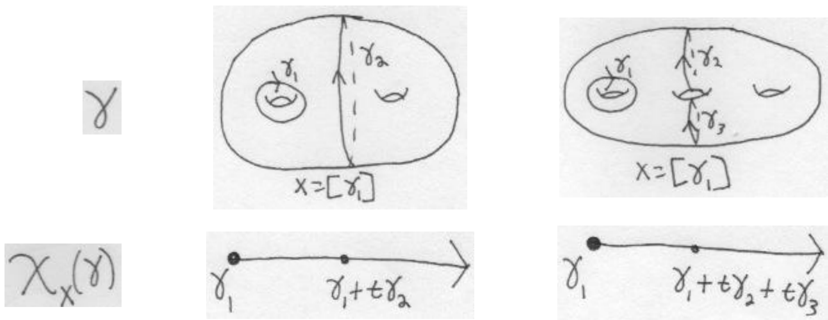

Figure 2: Examples of multicurves γ whose associated cells $\mathcal{X}_x(\gamma)$ are noncompact 1-dimensional polyhedra. Under each multicurve is how to write x as a linear combination of the multicurves.

3 Basic properties of cells

We now discuss some basic properties of cells. Throughout this section, $\gamma = \gamma_1 \cup \cdots \cup \gamma_k$ is a fixed oriented multicurve on Σ such that $\mathcal{X}_x(\gamma)$ is nondegenerate.

Zero sets. Define $\mathcal{Z}(\gamma)$ to be the set of all weighted oriented multicurves $c_1\gamma_1 + \cdots + c_k\gamma_k$ that represent $0 \in \mathrm{H}_1(\Sigma;\mathbb{Z})$. Just like for $\mathcal{X}_x(\gamma)$, we can identify $\mathcal{Z}(\gamma)$ with a subset (in fact, a linear subspace) of $\mathbb{R}^k$. Since we are assuming that $\mathcal{X}_x(\gamma)$ is nondegenerate, under these identifications $\mathcal{X}_x(\gamma)$ is the intersection of an affine subset of $\mathbb{R}^k$ parallel to $\mathcal{Z}(\gamma)$ with the positive orthant $\mathbb{R}^k_{\geq 0}$.

Generators and relations for zero sets. Let R be a subsurface of Σ whose boundary components (considered as *unoriented* curves) lie in γ. Letting $1 \leq i_1 < i_2 < \cdots < i_p \leq k$ be the indices such that the boundary components of R are $\gamma_{i_1}, \ldots, \gamma_{i_p}$, define

$$\partial R = \pm\gamma_{i_1} + \cdots + \pm\gamma_{i_p},$$

where the signs reflect whether or not the orientation of γ_{i_j} agrees or not with the orientation it acquires from R. Clearly $\partial R \in \mathcal{Z}(\gamma)$. We then have the following.

Lemma 3.1. *Let γ be an oriented multicurve on Σ. Let $R_1, \ldots, R_\ell$ be the connected subsurfaces of Σ obtained by cutting Σ along γ. Then $\mathcal{Z}(\gamma)$ is*

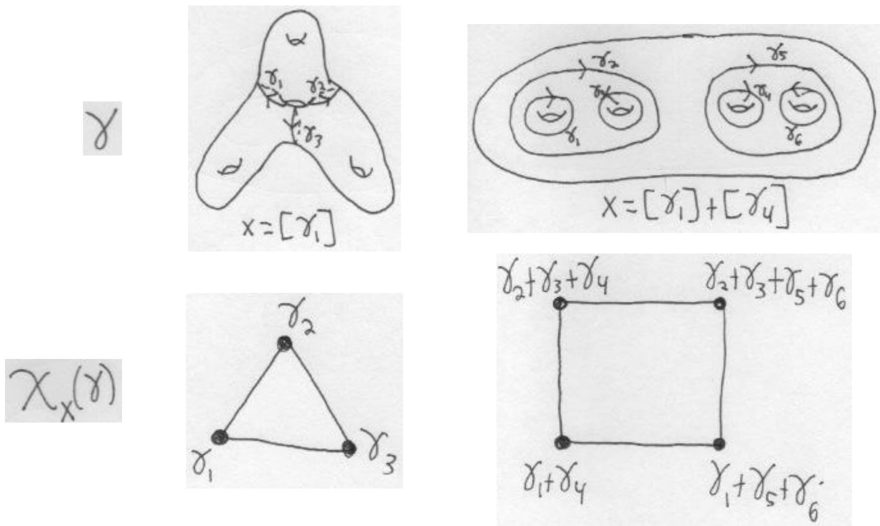

Figure 3: Examples of multicurves γ whose associated cells $\mathcal{X}_x(\gamma)$ are compact 2-dimensional polyhedra. Under each multicurve is how to write x as a linear combination of the multicurves. The points in the interior of the triangle are $t_1\gamma_1 + t_2\gamma_2 + t_3\gamma_3$ with $t_1, t_2, t_3 \geq 0$ and $t_1 + t_2 + t_3 = 1$. The points in the interior of the square are $t\gamma_1 + s\gamma_2 + s\gamma_3 + t'\gamma_4 + s'\gamma_5 + s'\gamma_6$ with $t, t', s, s' \geq 0$ and $t + s = 1$ and $t' + s' = 1$. We remark that one can also find cells that are pentagons, hexagons, etc.

generated by $\{\partial R_1, \ldots, \partial R_\ell\}$*, and the only relation between these generators is* $\partial R_1 + \cdots + \partial R_\ell = 0$.

Proof. We can clearly identify $\mathcal{Z}(\gamma)$ with the kernel of the map $\mathrm{H}_1(\gamma;\mathbb{Z}) \to \mathrm{H}_1(\Sigma;\mathbb{Z})$. The long exact sequence in homology associated to the pair (Σ, γ) therefore induces an exact sequence

$$\mathrm{H}_2(\Sigma;\mathbb{Z}) \longrightarrow \mathrm{H}_2(\Sigma/\gamma;\mathbb{Z}) \xrightarrow{\pi} \mathcal{Z}(\gamma) \longrightarrow 0.$$

Letting $\overline{R}_i$ be the image of R_i in Σ/γ (see Figure 4), we have an element $[\overline{R}_i] \in \mathrm{H}_2(\Sigma/\gamma;\mathbb{Z})$ satisfying $\pi([\overline{R}_i]) = \partial R_i$. The group $\mathrm{H}_2(\Sigma/\gamma;\mathbb{Z})$ is the free abelian group with basis $\{[\overline{R}_1], \ldots, [\overline{R}_\ell]\}$, and the generator $[\Sigma]$ of $\mathrm{H}_2(\Sigma;\mathbb{Z}) \cong \mathbb{Z}$ maps to $[\overline{R}_1] + \cdots + [\overline{R}_\ell]$. The lemma follows. □

Lemma 3.1 has the following immediate corollary.

Corollary 3.2. *Let* γ *be an oriented multicurve on* Σ *such that* $\mathcal{X}_x(\gamma)$ *is nondegenerate. Let* $\ell \geq 1$ *be the number of components of* Σ *cut along* γ*. Then* $\mathcal{X}_x(\gamma)$ *is an* $(\ell - 1)$*-dimensional (not necessarily compact) polyhedron.*

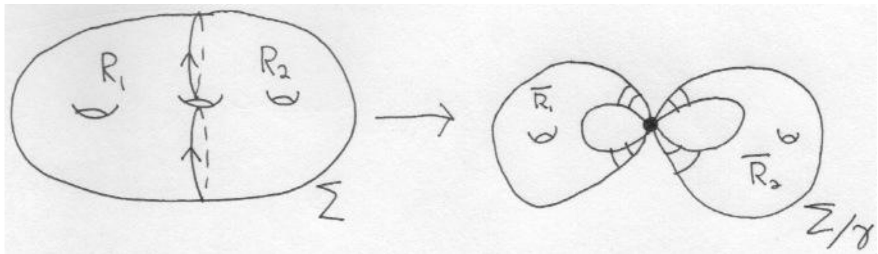

Figure 4: The result Σ/γ of collapsing a multicurve to a point. For $i = 1, 2$, the surface $\overline{R}_i$ can both be obtained from a torus by identifying two of its points together.

Vertices. We now give a concrete description of the vertices of $\widehat{\mathcal{C}}_x(\Sigma)$ (which of course coincide with the vertices of $\mathcal{C}_x(\Sigma)$).

Lemma 3.3. *The vertices of $\widehat{\mathcal{C}}_x(\Sigma)$ consist of $c_1\gamma_1 + \cdots + c_k\gamma_k$, where $\gamma_1 \cup \cdots \cup \gamma_k$ is an oriented multicurve on Σ that does not separate Σ and the c_i are positive integers such that*

$$c_1[\gamma_1] + \cdots + \cdots c_k[\gamma_k] = x.$$

Proof. Consider a point $c = c_1\gamma_1 + \cdots + c_k\gamma_k$ of $\widehat{\mathcal{C}}_x(\Sigma)$. Assume that none of the c_i vanish, so the support of c is $\gamma = \gamma_1 \cup \cdots \cup \gamma_k$. The cell $\mathcal{X}_x(\gamma)$ is therefore nondegenerate. The point c is a vertex of $\widehat{\mathcal{C}}_x(\Sigma)$ exactly when $\mathcal{X}_x(\gamma)$ is 0-dimensional (and hence consists of the single point c). Corollary 3.2 says that this holds if and only if $\gamma_1 \cup \cdots \cup \gamma_k$ does not separate Σ. It remains to prove that if c is a vertex, then each c_i is an integer.

Consider some $1 \leq i \leq k$. Since $\gamma_1 \cup \cdots \cup \gamma_k$ does not separate Σ, we can find an oriented simple closed curve δ on Σ that intersects γ_i once with a positive sign and is disjoint from γ_j for $1 \leq j \leq k$ with $j \neq i$. We then have

$$c_i = \hat{i}([\delta], c_1[\gamma_1] + \cdots + c_k[\gamma_k]) = \hat{i}([\delta], x) \in \mathbb{Z};$$

the final inequality follows from the fact that $x \in \mathrm{H}_1(\Sigma; \mathbb{Z})$. □

Criterion for being reduced. We now prove the following simple description of when a cell is reduced (compare with the examples in Figure 2).

Lemma 3.4. *Let $\gamma = \gamma_1 \cup \cdots \cup \gamma_k$ be a multicurve such that that $\mathcal{X}_x(\gamma)$ is nondegenerate. The cell $\mathcal{X}_x(\gamma)$ is reduced if and only if there does not exist $1 \leq i_1 < \cdots < i_p \leq k$ such that $[\gamma_{i_1}] + \cdots + [\gamma_{i_p}] = 0$.*

Proof. If there exist $1 \leq i_1 < \cdots < i_p \leq k$ such that $[\gamma_{i_1}] + \cdots + [\gamma_{i_p}] = 0$, then fixing some point $c \in \mathcal{X}_x(\gamma)$ we have an infinite ray

$$\{c + t(\gamma_{i_1} + \cdots + \gamma_{i_p}) \mid t \geq 0\} \subset \mathcal{X}_x(\gamma).$$

Thus $\mathcal{X}_x(\gamma)$ is noncompact, and hence not reduced.

We will prove the contrapositive of the other implication of the lemma. Assume that $\mathcal{X}_x(\gamma)$ is nonreduced, i.e. not compact. We first prove that there exist real numbers $c_1, \ldots, c_k \geq 0$ (not all 0) such that $c_1[\gamma_1] + \cdots + c_k[\gamma_k] = 0$. Since $\mathcal{X}_x(\gamma)$ is a noncompact polyhedron, it must contain an infinite ray. Let $c' = c_1'\gamma_1 + \cdots + c_k'\gamma_k$ be the initial point of this ray and let $c'' = c_1''\gamma_1 + \cdots + c_k''\gamma_k$ be some other point on this ray. For $1 \leq i \leq k$, set $c_i = c_i'' - c_i'$. We thus have

$$c_1[\gamma_1]+\cdots+c_k[\gamma_k] = (c_1'[\gamma_1]+\cdots+c_k'[\gamma_k])-(c_1''[\gamma_1]+\cdots+c_k''[\gamma_k]) = x-x = 0.$$

Moreover, the points

$$\{(c_1' + tc_1)\gamma_1 + \cdots + (c_k' + tc_k)\gamma_k \mid t \geq 0\}$$

all lie in $\mathcal{X}_x(\gamma)$, i.e. $c_i' + tc_i \geq 0$ for all $t \geq 0$ and all $1 \leq i \leq k$. We conclude that $c_i \geq 0$, as desired.

Let $R_1, \ldots, R_\ell$ be the connected subsurfaces of Σ obtained by cutting Σ along γ. Lemma 3.1 implies that there exists some $d_1, \ldots, d_\ell \in \mathbb{R}$ such that

$$c_1\gamma_1 + \cdots + c_k\gamma_k = d_1\partial R_1 + \cdots d_\ell\partial R_\ell.$$

Since $\partial R_1 + \cdots + \partial R_\ell = 0$, we can add a large positive constant E to each d_i and ensure that $d_i > 0$ for $1 \leq i \leq \ell$. Set $d = \max\{d_1, \ldots, d_\ell\}$, and assume that the R_j are ordered such that $d_1 = \cdots = d_r = d$ and $d_{r+1}, \ldots, d_\ell < d$. Since not all the c_i are 0 and $\partial R_1 + \cdots + \partial R_\ell = 0$, we must have $r < \ell$. Setting $R = R_1 \cup \cdots \cup R_r$, the surface R is thus a proper subsurface of Σ, so $\partial R \neq 0$. Observe that

$$c_1\gamma_1 + \cdots + c_k\gamma_k = d\partial R + d_{r+1}\partial R_{r+1} + \cdots + d_\ell\partial R_\ell.$$

Each γ_i occurs as the boundary of exactly two of the R_j. Since $d_j < d$ for $r + 1 \leq j \leq \ell$ and $c_i \geq 0$ for all $1 \leq i \leq k$, the coefficients of all of the γ_i which appear in ∂R must be $+1$ (as opposed to -1). In other words,

$$\partial R = \gamma_{i_1} + \cdots + \gamma_{i_p}$$

for some $1 \leq i_1 < \cdots < i_p \leq k$, as desired. $\square$

4 Prerequisites for contractibility

As we said at the beginning of this note, our main result will be that the complex $\mathcal{C}_x(\Sigma)$ is contractible. This will be proven in the next section after we discuss some preliminary results. The heart of our proof will be an explicit deformation retraction of $\widehat{\mathcal{C}}_x(\Sigma)$ to a point; we will then deduce that $\mathcal{C}_x(\Sigma)$ is contractible by giving an explicit (and fairly simple) deformation retraction of $\widehat{\mathcal{C}}_x(\Sigma)$ to $\mathcal{C}_x(\Sigma)$. To construct a deformation retraction of $\widehat{\mathcal{C}}_x(\Sigma)$ to a point, we will construct canonical "straight lines" between any two points in $\widehat{\mathcal{C}}_x(\Sigma)$. This will be done via a parameterization of $\widehat{\mathcal{C}}_x(\Sigma)$ by a set of differential forms; see the map Λ constructed below. While the entirety of this set of differential forms is not convex, it is close enough to being convex that we can use it to get the desired "straight lines".

Hyperbolic geometry. We will need tiny amount of hyperbolic geometry. Recall that a hyperbolic metric is a Riemannian metric with constant sectional curvature -1. These exist on all closed surfaces whose genus is at least 2. Fixing a hyperbolic metric on Σ, the following three facts then hold.

- Every nonnullhomotopic simple closed curve on Σ is homotopic to a unique simple geodesic.
- Any two distinct simple geodesics on Σ intersect transversely.
- Let γ and γ' be disjoint nonnullhomotopic simple closed curves on Σ. Assume that γ and γ' are not homotopic to each other. Then the geodesics that are homotopic to γ and γ' are disjoint.

See [3, §1.1] for this and much more.

Cleaning up curves. If $\gamma_1 \cup \cdots \cup \gamma_k$ is a collection of disjoint oriented simple closed curves on Σ and $c_1, \ldots, c_k \in \mathbb{R}$, then $c_1\gamma_1 + \cdots + c_k\gamma_k$ is not necessarily a positively weighted oriented multicurve : some of the c_i might be negative, some of the γ_i might be homotopic to each other (possibly with opposite orientations), and some of the γ_i might be nullhomotopic. However, by discarding the nullhomotopic γ_i, reversing the orientations of some of the γ_i (and changing the signs of the corresponding c_i), and collecting together the homotopic γ_i, we obtain a canonical positively weighted oriented multicurve c. We will say that c is obtained by *cleaning up* $c_1\gamma_1 + \cdots + c_k\gamma_k$. This definition extends in an obvious way if some of the γ_i are oriented 1-submanifolds with multiple components.

Maps to circle. Consider a smooth map $f : \Sigma \to S^1$. For any regular value $p \in S^1$ of f, the pullback $f^{-1}(p)$ is an oriented 1-submanifold of Σ. We will say that f *represents* the associated element $[f^{-1}(p)]$ of $\mathrm{H}_1(\Sigma; \mathbb{Z})$; this makes sense since $[f^{-1}(p)] = [f^{-1}(q)]$ for any two regular values $p, q \in S^1$. This latter assertion follows from the fact that if λ is an oriented arc of S^1 with oriented boundary $p - q$, then $f^{-1}(\lambda)$ is a subsurface whose oriented boundary is $f^{-1}(p) \sqcup -f^{-1}(q)$. Another way of describing $[f^{-1}(p)]$ is that it is the element of $\mathrm{H}_1(\Sigma; \mathbb{Z})$ which is Poincaré dual to $f^*([S^1]) \in \mathrm{H}^1(\Sigma; \mathbb{Z})$. This can be derived from the relationship between cup products on cohomology and intersection products on homology; see, e.g., [2, §VI.12].

Weighted multicurve from map to circle. Now assume that $f : \Sigma \to S^1$ represents $x \in \mathrm{H}_1(\Sigma; \mathbb{Z})$ and has finitely many critical values. These critical values divide S^1 into arcs $\lambda_1, \ldots, \lambda_n$. Normalize S^1 such that its circumference is 1. Setting $c_i = \text{length}(\lambda_i)$, we thus have $c_1 + \cdots + c_n = 1$. For all $1 \leq i \leq n$, let q_i be an arbitrary point in the interior of λ_i and let $\delta_i = f^{-1}(q_i)$. Since $[\delta_i] = x$ for all $1 \leq i \leq n$, we have

$$c_1[\delta_1] + \cdots + c_q[\delta_n] = (c_1 + \cdots + c_n)[x] = [x].$$

Define $\Psi(f) \in \widehat{\mathcal{C}}_x(\Sigma)$ to be the result of cleaning up $c_1\delta_1 + \cdots + c_n\delta_n$. The element $\Psi(f)$ appears to depend on the choice of the q_i; however, different choices of q_i will yield homotopic δ_i, and thus $\Psi(f)$ is well-defined.

Globalizing the construction. Let $\mathcal{F}_x(\Sigma, S^1)$ be the space of smooth maps $\Sigma \to S^1$ representing x which have finitely many critical values. Give $\mathcal{F}_x(\Sigma, S^1)$ the C^∞-topology. The above construction yields a map $\Psi : \mathcal{F}_x(\Sigma, S^1) \to \widehat{\mathcal{C}}_x(\Sigma)$.

Lemma 4.1. *The map $\Psi : \mathcal{F}_x(\Sigma, S^1) \to \widehat{\mathcal{C}}_x(\Sigma)$ is continuous.*

Proof. As f moves around $\mathcal{F}_x(\Sigma, S^1)$, the critical values of f move continuously around S^1. The 1-submanifolds of Σ used to define $\Psi(f)$ therefore also move homotopically around in Σ. When two critical values come together (causing one of the arcs used to define $\Psi(f)$ to disappear), the weight on the corresponding submanifold of Σ shrinks to 0. □

The following two lemmas show that Ψ is insensitive to certain deformations of its input.

Lemma 4.2. *Let $r : S^1 \to S^1$ be a rotation. Then for all $f \in \mathcal{F}_x(\Sigma, S^1)$ we have $\Psi(f) = \Psi(r \circ f)$.*

Proof. Obvious. □

Lemma 4.3. *Let $f_t \in \mathcal{F}_x(\Sigma, S^1)$ be a continuous family of maps for $t \in [0,1]$. Assume that the critical values of f_t and $f_{t'}$ are equal for all $t, t' \in [0,1]$. Then $\Psi(f_0) = \Psi(f_1)$.*

Proof. Let $\lambda_1, \ldots, \lambda_n$ be the arcs into which S^1 is divided by the common critical values of the f_t, and set $c_i = \text{length}(\lambda_i)$. For $1 \leq i \leq n$, let q_i be an arbitrary point in the interior of λ_i. Finally, for $0 \leq t \leq 1$ let $\gamma_i(t) = f_t^{-1}(q_i)$. The key observation is that the curve $\gamma_i(t)$ depends continuously on t, so $\gamma_i(0)$ is homotopic to $\gamma_i(1)$. Thus

$$c_1\gamma_1(0) + \cdots + c_n\gamma_n(0) = c_1\gamma_1(1) + \cdots + c_n\gamma_n(1),$$

and the lemma follows. □

Constructing maps using one-forms. To make the above results useful, we need a way of constructing elements of $\mathcal{F}_x(\Sigma, S^1)$. Consider a smooth closed 1-form ω on Σ which is Poincaré dual to $x \in \mathrm{H}_1(\Sigma; \mathbb{Z})$. In other words,

$$\int_h \omega = \hat{i}(x, h) \qquad (h \in \mathrm{H}_1(\Sigma; \mathbb{Z})).$$

For all basepoints $p_0 \in \Sigma$, we can define a smooth map $\Phi(\omega, p_0) : \Sigma \to S^1$ as follows. Regard S^1 as $\mathbb{R}/\mathbb{Z}$. The for any $q \in \Sigma$, we define $\Phi(\omega, p_0)(q)$ to be the image of $\int_\alpha \omega$ in S^1, where α is a smooth path on Σ from p_0 to q. The fact that ω is Poincaré dual to x implies that the integral of ω around any closed loop is an integer, so this is well-defined. The critical points of $\Phi(\omega, p_0)$ are exactly the zeros of ω.

Define $\Omega_x(\Sigma)$ to be the set of smooth closed 1-forms ω on Σ with the following two properties.

- ω is Poincaré dual to $x \in \mathrm{H}_1(\Sigma; \mathbb{Z})$.
- The zero set of ω has finitely many connected components.

Endow $\Omega_x(\Sigma)$ with the C^∞-topology. The above discussion is summarized in the following lemma.

Lemma 4.4. *There exists a continuous map $\Phi : \Omega_x(\Sigma) \times \Sigma \to \mathcal{F}_x(\Sigma, S^1)$.*

Proof. The only new assertion here is the continuity of Φ, but this is obvious from its definition. □

Changing the basepoint p_0 has the following effect on $\Phi(\omega, p_0)$.

Lemma 4.5. *Let $\omega \in \Omega_x(\Sigma)$ and $p_0, p_0' \in \Sigma$. Then $\Phi(\omega, p_0') = r \circ \Phi(\omega, p_0)$, where $r : S^1 \to S^1$ is a rotation.*

Proof. We can take r to be a rotation of S^1 by $\int_\alpha \omega$, where α is a smooth path on Σ from p_0' to p_0. $\square$

Combining the constructions. Define a map $\Lambda : \Omega_x(\Sigma) \to \widehat{\mathcal{C}}_x(\Sigma)$ by setting $\Lambda(\omega) = \Psi(\Phi(\omega, p_0))$, where $p_0 \in \Sigma$ is an arbitrary base point. Lemmas 4.2 and 4.5 show that $\Lambda(\omega)$ does not depend on the choice of p_0. The main properties of Λ are contained in the following two lemmas.

Lemma 4.6. *The map $\Lambda : \Omega_x(\Sigma) \to \widehat{\mathcal{C}}_x(\Sigma)$ is continuous.*

Proof. An immediate consequence of Lemmas 4.1 and 4.4. $\square$

Lemma 4.7. *For $0 \leq t \leq 1$, let $\omega_t \in \Omega_x(\Sigma)$ be a continuous family of 1-forms. Assume that we can find a set $\{p_0, \ldots, p_\ell\}$ of points on Σ with the following properties.*

- *For all $0 \leq t \leq 1$, the set $\{p_0, \ldots, p_\ell\}$ consists of exactly one point in each connected component of the zero set of ω_t.*
- *For all $1 \leq i \leq \ell$, there exists an arc α_i in Σ connecting p_0 to p_i such that $\int_{\alpha_i} \omega_t = \int_{\alpha_i} \omega_{t'}$ for all $0 \leq t, t' \leq 1$.*

Then $\Lambda(\omega_0) = \Lambda(\omega_1)$.

Proof. By construction, for all $0 \leq t \leq 1$ the critical values of $\Phi(\omega_t, p_0)$ are exactly the images in $S^1 = \mathbb{R}/\mathbb{Z}$ of the set $\{0, \int_{\alpha_1} \omega_t, \ldots, \int_{\alpha_\ell} \omega_t\}$. The lemma thus follows from Lemma 4.3. $\square$

Example I : single curve. We now give the first of two examples of the above techniques. Let γ be an oriented simple closed curve on Σ such that $[\gamma] = x$. We will construct some $\omega \in \Omega_x(\Sigma)$ such that $\Lambda(\omega) = \gamma$. Assume that we have fixed a hyperbolic metric on Σ. Homotoping γ, we can assume that it is a geodesic. Parameterize the annulus $\mathbb{A}$ in polar coordinates as $\{(r, \theta) \mid 1 \leq r \leq 3 \text{ and } 0 \leq \theta < 2\pi\}$. For $\epsilon > 0$, an *ϵ-strip map* around γ is an embedding $\iota : \mathbb{A} \hookrightarrow \Sigma$ with the following properties.

- The map ι takes the oriented "core" curve $\{(r, \theta) \mid r = 2,\ 0 \leq \theta < 2\pi\}$ of $\mathbb{A}$ to γ, parameterized at constant speed.

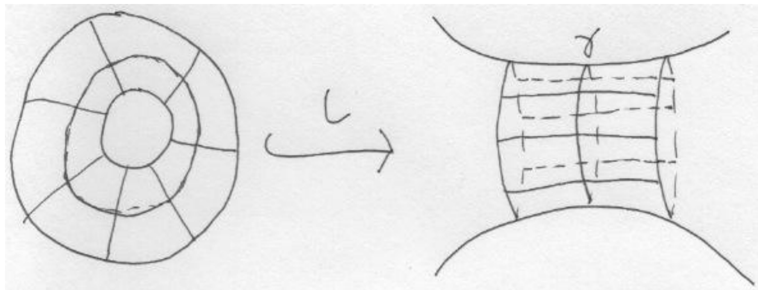

Figure 5: An ϵ-strip.

- For all angles $0 \leq \theta_0 < 2\pi$, the map ι takes the oriented line segment $\{(r, \theta_0) \mid 1 \leq r \leq 3\}$ in $\mathbb{A}$ to a geodesic segment of length 2ϵ that intersects γ orthogonally with a positive sign. Again, this geodesic segment is parameterized at constant speed.

See Figure 5. For $\epsilon > 0$ sufficiently small these exist and are unique up to precomposition with a rotation of $\mathbb{A}$. The image A of ι will be called an *ϵ-strip* around γ. Define $\mu : \mathbb{R} \to \mathbb{R}$ to be the function

$$\mu(x) = \begin{cases} \frac{1}{\int_{-\infty}^{\infty} e^{-1/(1-(z-2)^2)} \mathrm{d}z} \int_0^x e^{-1/(1-(z-2)^2)} \mathrm{d}z & \text{if } x \in [1,3], \\ 0 & \text{if } x \notin [1,3]. \end{cases}$$

Thus μ is a smooth nonnegative function of total integral 1 which is supported on $[1,3]$. There is a smooth closed 1-form $\mu(r)\mathrm{d}r$ on $\mathbb{A}$. We can therefore define a smooth closed 1-form ω on Σ via the formulas

$$\omega|_A = \iota_*(\mu(r)\mathrm{d}r) \quad \text{and} \quad \omega|_{\Sigma \setminus A} = 0.$$

We will call ω the *ϵ-strip form* dual to γ. It is clear that ω represents $[\gamma] = x$. Additionally, we have the following lemma

Lemma 4.8. *With the notation as above, we have $\Lambda(\omega) = \gamma$.*

Proof. Fix a basepoint $p_0 \in \Sigma \setminus A$. Regarding S^1 as $\mathbb{R}/\mathbb{Z}$, it is then clear from the definitions that $\Phi(\omega, p_0) : \Sigma \to S^1$ is the map

$$\Phi(\omega, p_0)(q) = \begin{cases} \int_1^r \mu(r)\mathrm{d}r & \text{if } q = \iota(r,\theta) \in A \text{ with } (r,\theta) \in \mathbb{A}, \\ 0 & \text{if } q \notin A. \end{cases}$$

In particular, the only critical value of $\Phi(\omega, p_0)$ is 0, and the preimage under $\Phi(\omega, p_0)$ of a regular value $q_1 \in (0,1) \subset S^1$ is a loop of the form

$\{\iota(r_1, \theta) \mid 0 \leq \theta < 2\pi\}$ for some $1 < r_1 < 3$. This loop is homotopic to γ, so we conclude that

$$\Lambda(\omega) = \Psi(\Phi(\omega, p_0)) = 1 \cdot \lambda = \lambda,$$

as desired. $\square$

Example II : multicurve. We now generalize the previous example. Let $c = c_1\gamma_1 + \cdots + c_k\gamma_k$ be an arbitrary positively weighted oriented multicurve on Σ which represents x. Again assume that we have fixed a hyperbolic metric on Σ and that each γ_i is a geodesic. Let $\epsilon > 0$ be small enough that there are ϵ-strips around each γ_i which are pairwise disjoint. For $1 \leq i \leq k$ let ω_i be the ϵ-strip form dual to γ_i, so ω_i represents $[\gamma_i]$. Finally, define $\omega = c_1\omega_1 + \cdots + c_k\omega_k$. It is then an easy exercise in the definitions to see that ω represents $c_1[\gamma_1] + \cdots + c_k[\gamma_k]$. Moreover, we have the following generalization of Lemma 4.8.

Lemma 4.9. *With the notation as above, we have $\Lambda(\omega) = c$.*

Proof. For $1 \leq i \leq k$, let A_i be the ϵ-strip around γ_i. Pick a basepoint $p_0 \in \Sigma \setminus \cup_{i=1}^k A_i$. Just like in the proof of Lemma 4.8, the map $\Phi(\Sigma, p_0) : \Sigma \to S^1$ takes p_0 to $0 \in S^1 = \mathbb{R}/\mathbb{Z}$, takes each component of $\Sigma \setminus \cup_{i=1}^k A_i$ to a critical value, and takes A_i to an arc of $S^1 = \mathbb{R}/\mathbb{Z}$ of length c_i (starting and ending at a critical value; observe that this arc can contain critical values in its interior). Let $\lambda_1, \ldots, \lambda_n$ be the arcs into which S^1 is divided by the critical values, let $d_i = \text{length}(\lambda_i)$, and let q_i be an arbitrary point in the interior of λ_i. Define $\delta_i = \Phi(\Sigma, p_0)^{-1}(q_i)$, so $\Lambda(\omega)$ is the result of cleaning up

$$d_1\delta_1 + \cdots + d_n\delta_n. \tag{1}$$

It is clear that $\cup_{j=1}^n \delta_j \subset \cup_{i=1}^k A_i$.

Fix some $1 \leq i \leq k$. The components of $\delta_1 \cup \cdots \cup \delta_n$ lying in A_i consist of a set of curves each of which is homotopic to γ_i. From (1), each of these curves has a weight from among the numbers $d_1, \ldots, d_n$. It is easy to see that these weights add up to c_i. The lemma follows. $\square$

5 Contractibility

We finally prove the following theorem of Bestvina–Bux–Margalit [1]. In its proof, we will use all of the notation introduced in §4.

Theorem 5.1. *Let Σ be a closed surface and let $x \in \mathrm{H}_1(\Sigma; \mathbb{Z})$ be a primitive vector. Then $\mathcal{C}_x(\Sigma)$ is contractible.*

Proof. The theorem has no content if the genus of Σ is 0 since in that case $\mathrm{H}_1(\Sigma;\mathbb{Z}) = 0$ contains no primitive vectors. If the genus of Σ is 1, then Σ contains no oriented multicurves with more than one component. The complex $\mathcal{C}_x(\Sigma)$ therefore is a discrete set of points, one for each homotopy class of oriented simple closed curve γ with $[\gamma] = x$. It is standard that in genus 1 there exists a unique such homotopy class of curves. We conclude that if the genus of Σ is 1, then $\mathcal{C}_x(\Sigma)$ consists of exactly one point and is hence contractible.

We can therefore assume without loss of generality that Σ has genus at least 2, which allows us to fix a hyperbolic metric on Σ. The proof of the theorem now has two steps.

Step 1. *The space $\widehat{\mathcal{C}}_x(\Sigma)$ is contractible.*

Let γ_0 be an oriented simple closed curve on Σ such that $[\gamma_0] = x$. Homotoping γ_0, we can assume that it is a hyperbolic geodesic. We will construct an explicit homotopy $f_t : \widehat{\mathcal{C}}_x(\Sigma) \to \widehat{\mathcal{C}}_x(\Sigma)$ such that $f_0 = \mathrm{id}$ and such that $f_1(c) = \gamma_0$ for all $c \in \widehat{\mathcal{C}}_x(\Sigma)$. This construction is divided into three substeps. In the first, we construct f_t on a fixed cell $\mathcal{X}_x(\gamma)$. This construction depends on a parameter $\epsilon > 0$; the second substep shows that in fact its output is independent of ϵ. The final substep shows how to piece together the maps on the various cells to define f_t.

Substep 1. *Let γ be an oriented multicurve such that $\mathcal{X}_x(\gamma)$ is nondegenerate. For all $\epsilon > 0$ sufficiently small, we construct a homotopy $f^{\epsilon}_{\gamma,t} : \mathcal{X}_x(\gamma) \to \widehat{\mathcal{C}}_x(\Sigma)$ such that $f^{\epsilon}_{\gamma,0}$ is the inclusion and $f^{\epsilon}_{\gamma,1}(c) = \gamma_0$ for all $c \in \mathcal{X}_x(\gamma)$.*

Write $\gamma = \gamma_1 \cup \cdots \cup \gamma_k$. Homotoping the γ_i, we can assume that they are all hyperbolic geodesics. For $0 \leq i \leq k$, let A^{ϵ}_i be an ϵ-strip around γ_i. Choosing $\epsilon > 0$ small enough, we can assume that that the following hold.

- For $1 \leq i < j \leq k$, we have $A^{\epsilon}_i \cap A^{\epsilon}_j = \emptyset$.
- For $1 \leq i \leq k$, the ϵ-strips A^{ϵ}_0 and A^{ϵ}_i intersect transversely as in Figure 6.

For $0 \leq i \leq k$, let ω^{ϵ}_i be the ϵ-strip form dual to γ_i. For a point $c_1\gamma_1 + \cdots + c_k\gamma_k$ of $\mathcal{X}_x(\gamma)$ and some $0 \leq t \leq 1$, the 1-form

$$t\omega^{\epsilon}_0 + (1-t)c_1\omega^{\epsilon}_1 + \cdots + (1-t)c_k\omega^{\epsilon}_k \tag{2}$$

represents x. Moreover, the following hold.

- For $t = 0$, the zero set of (2) is the complement of $A^{\epsilon}_1 \cup \cdots \cup A^{\epsilon}_k$.

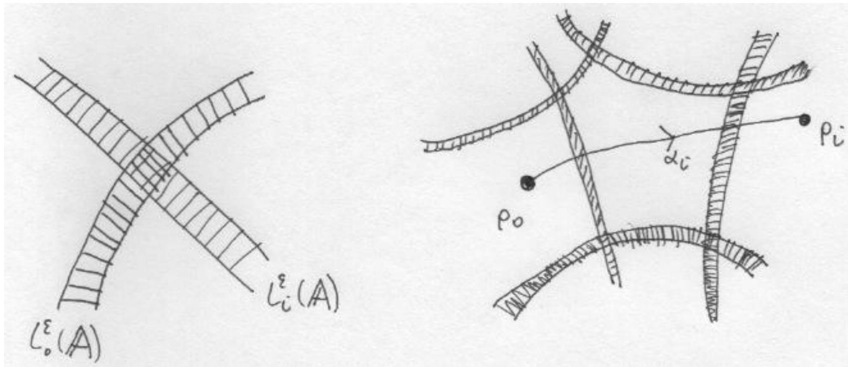

Figure 6: Left : Two transverse ϵ-strips. Right : The arc α_i crosses some of the ϵ-strips.

- For $0 < t < 1$, the zero set of (2) is the complement of $A_0^\epsilon \cup \cdots \cup A_k^\epsilon$ (this follows from our assumptions on the intersections of A_0^ϵ and A_i^ϵ for $1 \leq i \leq k$).
- For $t = 1$, the zero set of (2) is the complement of A_0^ϵ.

In particular, the zero set of (2) has finitely many components. The upshot of all of this is that (2) is an element of $\Omega_x(\Sigma)$ for all $0 \leq t \leq 1$. We can therefore define a function $f_{\gamma,t}^\epsilon : \mathcal{X}_x(\gamma) \to \widehat{\mathcal{C}}_x(\Sigma)$ via the formula

$$f_{\gamma,t}^\epsilon(c_1\gamma_1 + \cdots + c_k\gamma_k) = \Lambda(t\omega_0^\epsilon + (1-t)c_1\omega_1^\epsilon + \cdots + (1-t)c_k\omega_k^\epsilon).$$

Lemma 4.6 implies that $f_{\gamma,t}^\epsilon$ is continuous (both as a function and as a homotopy). Also, it follows from Lemma 4.9 that $f_{\gamma,0}^\epsilon$ is the inclusion and $f_{\gamma,1}^\epsilon(c) = \gamma_0$ for all $c \in \mathcal{X}_x(\gamma)$.

Substep 2. *Let γ be an oriented multicurve such that $\mathcal{X}_x(\gamma)$ is nondegenerate and let $\epsilon, \epsilon' > 0$ be small enough that $f_{\gamma,t}^\epsilon$ and $f_{\gamma,t}^{\epsilon'}$ are defined. Then $f_{\gamma,t}^\epsilon = f_{\gamma,t}^{\epsilon'}$.*

Without loss of generality, $\epsilon' < \epsilon$. As in Substep 1, write $\gamma = \gamma_1 \cup \cdots \cup \gamma_k$ with γ_i a hyperbolic geodesic for $1 \leq i \leq k$. Fix some $c_1\gamma_1 + \cdots + c_k\gamma_k \in \mathcal{X}_x(\gamma)$ and some $0 \leq t_0 \leq 1$. Our goal is to show that

$$f_{\gamma,t_0}^\epsilon(c_1\gamma_1 + \cdots + c_k\gamma_k) = f_{\gamma,t_0}^{\epsilon'}(c_1\gamma_1 + \cdots + c_k\gamma_k). \tag{3}$$

To simplify our notation, we will deal with the case where $0 < t_0 < 1$; the other cases are similar.

We now set up some notation. For $0 \leq i \leq k$ let A_i^ϵ be an ϵ-strip around γ_i. Also, for $0 \leq i \leq k$ and $\epsilon' \leq e \leq \epsilon$ let ω_i^e be the e-strip form dual to γ_i. Finally, for $\epsilon' \leq e \leq \epsilon$ let

$$\omega^e = t\omega_0^e + (1-t)c_1\omega_1^e + \cdots + (1-t)c_k\omega_k^e.$$

The assertion of (3) is thus equivalent to the assertion that $\Lambda(\omega^\epsilon) = \Lambda(\omega^{\epsilon'})$.

We will prove this using Lemma 4.7, whose conditions we now verify. First, by construction the differential forms ω^e depend continuously on e. Let $\{p_0, \ldots, p_\ell\}$ be a set of points on Σ that contains exactly one point in the interior of each component of

$$\Sigma \setminus \bigcup_{i=0}^{k} A_i^\epsilon.$$

Clearly $\{p_0, \ldots, p_\ell\}$ also contains exactly one point in the interior of each component of

$$\Sigma \setminus \bigcup_{i=0}^{k} A_i^e$$

for each $\epsilon' \leq e \leq \epsilon$. As we said in Substep 1, these are exactly the components of the zero set of ω^e (this is where we use the fact that $0 < t_0 < 1$). Finally, for $1 \leq i \leq \ell$ let α_i be any smooth arc from p_0 to p_i that crosses the γ_i transversely. Letting $\hat{i}(\alpha_i, \gamma_j)$ be the algebraic intersection number between the arc α_i and the simple closed curve γ_j, it is clear that for $1 \leq i \leq \ell$ and $\epsilon' \leq e \leq \epsilon$ we have

$$\int_{\alpha_i} \omega^e = t\hat{i}(\alpha_i, \gamma_0) + \sum_{j=1}^{k} (1-t)c_i\hat{i}(\alpha_i, \gamma_j).$$

See Figure 6. As this does not depend on e, the conditions of Lemma 4.7 are satisfied and we conclude that $\Lambda(\omega^\epsilon) = \Lambda(\omega^{\epsilon'})$, as desired.

Substep 3. *We construct a homotopy* $f_t : \widehat{\mathcal{C}}_x(\Sigma) \to \widehat{\mathcal{C}}_x(\Sigma)$ *such that* $f_0 = id$ *and such that* $f_1(c) = \delta_0$ *for all* $c \in \widehat{\mathcal{C}}_x(\Sigma)$.

If γ is any oriented multicurve such that $\mathcal{X}_x(\gamma)$ is nondegenerate, then using Substep 2 we can write $f_{\gamma,t} : \mathcal{X}_x(\gamma) \to \widehat{\mathcal{C}}_x(\Sigma)$ for $f_{\gamma,t}^\epsilon$, where $\epsilon > 0$ is an sufficiently small number. To show that the $f_{\gamma,t}$ piece together to give a function $f_t : \widehat{\mathcal{C}}_x(\Sigma) \to \widehat{\mathcal{C}}_x(\Sigma)$, it is enough to show that if γ and γ' are any oriented multicurves such that $\mathcal{X}_x(\gamma)$ and $\mathcal{X}_x(\gamma')$ are nondegenerate, then

$f_{\gamma,t}$ and $f_{\gamma',t}$ agree on the intersection of $\mathcal{X}_x(\gamma)$ and $\mathcal{X}_x(\gamma')$ in $\widehat{\mathcal{C}}_x(\Sigma)$. If this intersection is nonempty, then it is exactly $\mathcal{X}_x(\gamma'')$, where γ'' is the oriented multicurve consisting of all oriented simple closed curves that appear in both γ and γ'. But it is clear from their definitions that if $\epsilon > 0$ is small enough that all three of $f^{\epsilon}_{\gamma,t}$ and $f^{\epsilon}_{\gamma',t}$ and $f^{\epsilon}_{\gamma'',t}$ are defined, then all three of them agree on $\mathcal{X}_x(\gamma'')$.

Step 2. *The space $\widehat{\mathcal{C}}_x(\Sigma)$ deformation retracts to $\mathcal{C}_x(\Sigma) \subset \widehat{\mathcal{C}}_x(\Sigma)$.*

Consider a point $c = c_1\gamma_1 + \cdots + c_k\gamma_k$ in $\widehat{\mathcal{C}}_x(\Sigma)$. Discarding some the the γ_i, we can assume that $c_i > 0$ for all $1 \leq i \leq k$. We will write down a canonical (i.e. independent of all choices) path from c to $\mathcal{C}_x(\Sigma)$. It will be clear that this path depends continuously on c and that it is the constant path if $c \in \mathcal{C}_x(\Sigma)$.

If $c \notin \mathcal{C}_x(\Sigma)$, then the cell $\mathcal{X}_x(\gamma)$ is not reduced. Lemma 3.4 therefore implies that there exists some subsurface R of Σ such that

$$\partial R = \gamma_{i_1} + \cdots + \gamma_{i_p}$$

for some $1 \leq i_1 < \cdots < i_p \leq k$. Let $R_1, \ldots, R_q$ be *all* such subsurfaces. It follows that

$$\partial R_1 + \cdots + \partial R_q = d_1\gamma_1 + \cdots + d_k\gamma_k$$

for some $d_i \geq 0$ (not all 0). Setting $T = \min\{c_i/d_i \mid d_i > 0\}$, we have a path

$$t \mapsto c - t(\partial R_1 + \cdots + \partial R_q)$$

in $\widehat{\mathcal{C}}_x(\Sigma)$ defined for $0 \leq t \leq T$. At the endpoint of this path, the coefficient of at least one of the γ_i has become 0. Repeat this process until c ends up in $\mathcal{C}_x(\Sigma)$. □

References

[1] M. Bestvina, K.-U. Bux and D. Margalit, The dimension of the Torelli group, J. Amer. Math. Soc. **23** (2010), no. 1, 61–105.

[2] G. E. Bredon, *Topology and geometry*, Graduate Texts in Mathematics, 139, Springer, New York, 1997.

[3] B. Farb and D. Margalit, *A primer on mapping class groups*, Princeton Mathematical Series, 49, Princeton Univ. Press, Princeton, NJ, 2012.

[4] A. Hatcher and D. Margalit, Generating the Torelli group, Enseign. Math. (2) **58** (2012), no. 1-2, 165–188.

[5] G. Mess, The Torelli groups for genus 2 and 3 surfaces, Topology **31** (1992), no. 4, 775–790.

Andrew Putman
Department of Mathematics
Rice University, MS 136
6100 Main St.
Houston, TX 77005
andyp@math.rice.edu