GEOMETRY OF THE HOMOLOGY CURVE COMPLEX

INGRID IRMER

Abstract. Suppose S is a closed, oriented surface of genus $g \geq 2$. This paper investigates the geometry of the "homology multicurve complex", $\mathcal{HC}(S,\alpha)$, of S; a complex closely related to complexes studied by Bestvina-Bux-Margalit and Hatcher. A path in $\mathcal{HC}(S,\alpha)$ corresponds to a homotopy class of immersed surfaces in $S \times I$. This observation is used to devise a simple algorithm for constructing quasi-geodesics connecting any two vertices in $\mathcal{HC}(S,\alpha)$. It is proven that for $g \geq 3$ the best possible bound on the distance between two vertices in $\mathcal{HC}(S,\alpha)$ depends linearly on their intersection number, in contrast to the logarithmic bound obtained in the standard curve complex. For $g \geq 4$ it is shown that $\mathcal{HC}(S,\alpha)$ is not δ-hyperbolic.

1. Introduction

Suppose S is a closed oriented surface with genus $g \geq 2$. A *curve* c in S is a piecewise smooth, injective map of S^1 into S that is not null homotopic. A *multicurve* is a union of pairwise disjoint curves on S. When convenient, a curve is confused with its image in S.

Fix a nontrivial element α of $H_1(S,\mathbb{Z})$. The *homology curve complex*, $\mathcal{HC}(S,\alpha)$, is a simplicial complex whose vertex set is the set of all isotopy classes of oriented multicurves in S in the homology class α. A set of vertices $m_1 ... m_k$ spans a simplex if the representatives of the isotopy classes can all be chosen to be disjoint.

The *distance*, $d_{\mathcal{H}}(v_1, v_2)$, between two vertices v_1 and v_2 is defined to be the distance in the path metric of the one-skeleton, where all edges have length one.

The *Torelli group* is the subgroup of the mapping class group that acts trivially on homology. $\mathcal{HC}(S,\alpha)$ is closely related to a complex defined in [1] that was used for calculating cohomological properties of the Torelli group.

Metric properties of curve complexes have been used for example for studying mapping class groups and the structure of 3-manifolds, for example [4], [14] and [9]. The aim of this paper is to study some basic geometric properties of $\mathcal{HC}(S,\alpha)$.

In [13] and [2] it was shown that the standard curve complex, $\mathcal{C}(S)$, is δ-hyperbolic. In contrast, in section 6 it will be shown that

Theorem. *For $g > 3$ and $\alpha \neq 0$, $\mathcal{HC}(S,\alpha)$ is not δ-hyperbolic.*

It is also well known (for example [12]) that in $\mathcal{C}(S)$, the distance between two vertices representing the curves a and b is less than or equal to $\log_2(i(a,b)) + 1$. However, in section 6 it will be shown that

Theorem. *$d_{\mathcal{H}}(m_1, m_2) \leq \frac{i(m_1,m_2)}{2} + 1$, where $i(m_1, m_2)$ is the geometric intersection number. This bound is sharp.*

An edge in $\mathcal{HC}(S,\alpha)$ connecting two vertices representing the multicurves γ_i and γ_{i+1} is called *simple* if $\gamma_{1+1} - \gamma_i$ is the oriented boundary of an embedded subsurface of S. A *simple path* is a path that only traverses simple edges. In proposition 3 an algorithm for constructing simple paths (hereafter referred to as the "path construction algorithm") is given.

Let I be a closed interval. In section 2.1 a path in $\mathcal{HC}(S,\alpha)$ connecting the vertices representing m_1 and m_2 is shown to correspond to an oriented, immersed surface H in $S \times I$ with ∂H homotopic to the multicurves $m_2 - m_1$ in $S \times 0$. The geometry of $\mathcal{HC}(S,\alpha)$ is thus related to the topology of surfaces in $S \times I$. In a later paper it will be shown that, modulo a uniformly bounded multiplicative constant, the distance between two vertices in $\mathcal{HC}(S,\alpha)$ representing the multicurves m_1 and m_2 is equal to the smallest possible genus of an orientable surface in $S \times I$ with boundary $m_2 - m_1$. In order to show that the path construction algorithm is optimal in some sense, the geometry of $\mathcal{HC}(S,\alpha)$ is related to the topology of immersed surfaces in $S \times I$ by defining two functions from $S \setminus (m_1 \cup m_2) \to \mathbb{Z}$: the "overlap" and the "pre-image function". These functions will now be briefly described.

Intersection numbers. There are two types of intersection numbers used in this work. The intersection number, also known as the *geometric intersection number*, is denoted by $i(m_1, m_2)$, and the *algebraic intersection number* by $\hat{i}(m_1, m_2)$. The algebraic intersection number of an oriented arc a with an oriented representative m_1 of the isotopy class $[m_1]$ is also written as $\hat{i}(a, m_1)$. Multicurves in $S \times I$ and their intersections with other multicurves are defined by projecting onto $S \times 0$.

Let π be the projection of $S \times I$ onto $S \times 0$ given by $(s, r) \mapsto s \times 0$. Informally, given an oriented, immersed surface H in $S \times I$, the *pre-image function*, $g_H : S \times 0 \setminus \pi(\partial H) \to \mathbb{Z}$ is given by $g_H(s) = \hat{i}(\pi^{-1}(s), H)$ (See Section 4 for a more precise definition). It is shown that, modulo an additive constant, the pre-image function does not depend on H but only on its boundary (lemma 4).

The overlap function, and the homological distance. The *overlap* f of two homologous multicurves m_1 and m_2 is the locally constant, upper semicontinuous function defined on S with minimum value zero and such that, for any two points x and y in $S\setminus(m_1 \cup m_2)$, $f(x)-f(y)$ is the algebraic intersection number of m_2-m_1 with an oriented arc with starting point y and endpoint x. The overlap of a null homologous multicurve n with itself is defined analogously. f is not dependent on the choice of oriented arc, because the algebraic intersection number of any closed loop with $m_2 - m_1$ is zero. The overlap does however depend on the choice of representatives of the isotopy classes $[m_1]$ and $[m_2]$. It will be assumed that the representatives of the free homotopy classes are chosen so that the maximum, M, of f is as small as possible. M will be called the homological distance, $\delta(m_1, m_2)$, between m_1 and m_2.

If H is a surface constructed from a simple path connecting m_1 and m_2, as described in subsection 2.1, the relation between g_H and the overlap of m_1 and m_2 shown in lemma 6 is used to show that the path construction algorithm constructs the shortest possible simple paths.

Theorem 7. Let m_1 and m_2 be two multicurves corresponding to vertices of $\mathcal{HC}(S, \alpha)$. Then the shortest simple path connecting the vertices has length equal to $\delta(m_1, m_2)$. Recall that by definition α is non trivial.

The path construction algorithm is similar to a construction in [8] for showing contractibility of the cyclic cycle complex, and can also be used to construct paths in this complex. It will be shown in corollary 8 that the paths so constructed in the cyclic cycle complex are geodesics.

A nice property of the path construction algorithm is that, as shown in theorem 3, it constructs the same unoriented path from m_1 to m_2 as from m_2 to m_1.

One reason for being interested in simple paths is that they are a simple means of estimating distance.

Theorem 9. If m_1 and m_2 do not contain null homologous submulticurves or homotopic curves, $d(m_1, m_2) < -3\chi(S)\delta(m_1, m_2)$.

1.1. **The Case $\alpha = 0$.** The case in which α is allowed to be null homologous is quite different. For example, in this case the complex admits an action of the full mapping class group, and when alpha is nontrivial, it does not. In the latter case, the natural group that acts is the subgroup of the mapping class group preserving alpha. Various complexes of null homologous (multi)curves, have been studied, for example the complex of separating curves and the Torelli geometry. Some of the methods discussed in this paper generalise, however the main problem seems to be that performing surgeries on null homologous multicurves could give trivial curves.

Acknowledgements. I would like to thank Ursula Hamenstädt for her supervision of this project. Also, without the advice and enthusiasm of many people, the writing of this paper could have dragged on into infinity. Thanks to Joan Birman, Benson Farb, Andrew Putman and Kasra Rafi. I am particularly grateful to Dan Margalit for his patience in teaching me how to write papers, and to Alan Hatcher for his detailed comments and improvements.

2. Simple Paths

The notion of a "simple path" is introduced in order to be able to perform counting arguments that relate surfaces in $S \times I$ to paths in $\mathcal{HC}(S, \alpha)$.

If n bounds an embedded subsurface of S, the union of the components of $S \setminus n$ whose boundary orientation coincides with the orientation of n will be called the *subsurface of S bounded by n.*

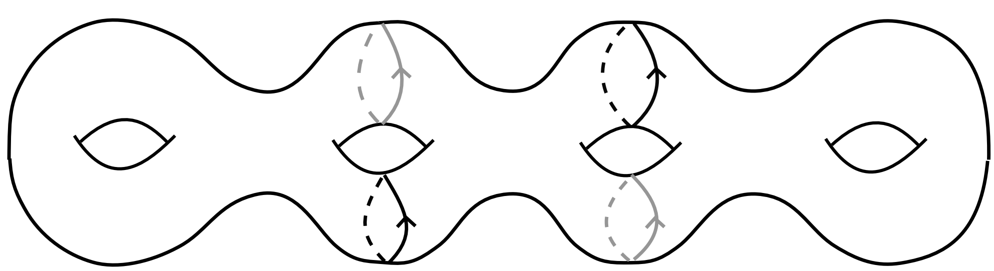

Figure 1. An edge that is not simple. The multicurve drawn in grey represents one vertex and the multicurve drawn in black represents the other.

The next lemma is used to decompose null homologous multicurves up into boundaries of subsurfaces, i.e. given a path in $\mathcal{HC}(S,\alpha)$, a simple path can easily be obtained by adding more edges where necessary.

Lemma 1. *If a null homologous multicurve n does not contain a nontrivial null homologous submulticurve, it bounds a subsurface of S.*

Proof. Consider the subsurface of S on which the overlap function of n has its maximum. This boundary is a null homologous submulticurve of n. By assumption on n it must be all of n.

□

Corollary 2. *Every path in $\mathcal{HC}(S,\alpha)$ can be made simple by adding extra vertices where necessary.*

Proof. Suppose m_1 and m_2 are connected by an edge that is not simple. By the previous lemma, $m_2 - m_1$ can be decomposed into k null homologous submulticurves $n_1, n_2, ...n_k$, each of which bounds a aubsurface of S. Then a simple path connecting m_1 and m_2 is determined by the vertices $m_1, m_1 - n_1, m_1 - n_1 - n_2, ...m_1 - n_1 - n_2 - ... - n_{k-1}, m_2$. □

2.1. **Constructing an Immersed Surface in $S \times I$ from a Path in $\mathcal{HC}(S,\alpha)$.** All curves, surfaces, and manifolds discussed here are assumed to be piecewise smooth.

Suppose γ is a simple path in $\mathcal{HC}(S,\alpha)$ passing through the vertices corresponding to the multicurves $\gamma_0, \gamma_1, ...\gamma_j$. A surface T_γ contained in $S \times j$ is constructed inductively. Given γ_0, isotope γ_1 such that there is a subsurface S_1 of S with boundary $\gamma_1 - \gamma_0$. Let T_1 be the surface in $S \times [0,1]$ given by $\gamma_0 \times [0, \frac{1}{2}] \cup S_1 \times \{\frac{1}{2}\} \cup \gamma_1 \times [\frac{1}{2}, 1]$. Next, isotope γ_2 so that there is a subsurface S_2 of S with $\partial S_2 = \gamma_2 - \gamma_1$ and let $T_2 = \gamma_1 \times [1, \frac{3}{2}] \cup S_2 \times \{\frac{3}{2}\} \cup \gamma_2 \times [\frac{3}{2}, 2]$. Repeat this successively for each of the γ_i until an embedded surface $T_\gamma = T_1 \cup T_2 \cup ... \cup T_j$ in $S \times [0, j]$ is obtained.

T_γ is called the *trace* of the path γ. Note that the trace of a path depends on the orientation on S.

Remark Similarly, if $\gamma_0, \gamma_1, ...\gamma_j$ is not simple, it can be used to construct a cell complex with boundary $\gamma_j - \gamma_0$. It is not difficult to show that such cell complexes are homotopic to immersed surfaces in $S \times [0, j]$.

2.2. **Extrema of f.** In order to construct paths in $\mathcal{HC}(S,\alpha)$, it is necessary to use some properties of the level sets, in particular the

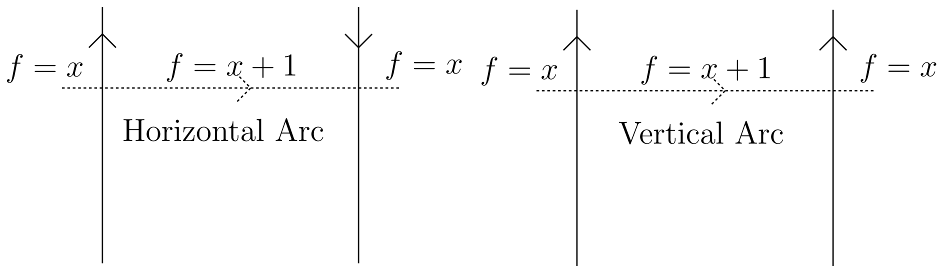

FIGURE 2. When the overlap f is thought of as a height function, a horizontal arc is horizontal with respect to f, and a vertical arc vertical.

local extrema, of f. These are used to define the surgeries used in the path construction algorithm.

Given an oriented multicurve a with a regular neighbourhood $\mathcal{N}(a)$ and an orientation on S, the left and right component of $\mathcal{N} \setminus a$ can be defined. If b is an oriented multicurve that intersects a transversely at a point p, it therefore makes sense to say that b *crosses over* a *from left to right* (or right to left) at p. Similarly, if b is an oriented arc with an endpoint on a, a notion in which b *leaves or approaches* a *from the left or right* can be defined.

Given two multicurves a and b on an oriented surface S, a *horizontal arc* of a is a component of $a \cap (S \setminus b)$ that leaves and approaches b from the same side. A *vertical arc* of $a \cap (S \setminus b)$ leaves and approaches b from opposite sides. An "innermost" arc in [8] is an example of a horizontal arc. Horizontal arcs are used to perform surgeries.

If a horizontal arc of $a \cap (S \setminus b)$ leaves and approaches b from the right, then this arc is *to the right of* b and vice versa.

Suppose a and b are multicurves in S in general position. Two arcs of $a \cap (S \setminus b)$ will be called homotopic if they are homotopic relative to b. Two oriented arcs will be said to be homotopic and oriented in the same way if one can be homotoped into the other in such a way that the orientations coincide.

It is not difficult to see that "verticalness" and "horizontalness" are properties of homotopy classes of arcs. Also, if $S \setminus (a \cup b)$ does not contain any bigons, a horizontal arc of $a \cap (S \setminus b)$ to the right of b can't be homotopic to a horizontal arc of $a \cap (S \setminus b)$ to the left of b, and an oriented arc of $a \cap (S \setminus b)$ is not homotopic to itself with the opposite orientation.

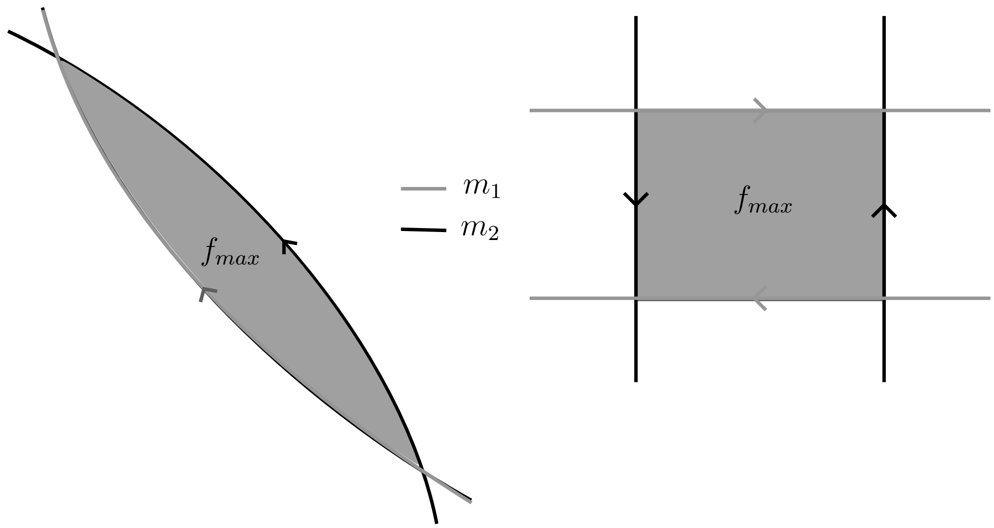

FIGURE 3. Examples of f_{max}.

Given m_1 and m_2, the overlap is bounded and has a maximum. Call the subsurface of S on which f takes on its maximum f_{max}. f_{max} has at least one connected component. The boundary of f_{max} consists of arcs of m_1 and m_2 such that f_{max} is to the right of any arc of m_1 on its boundary and to the left of any arc of m_2 on its boundary. In other words, the boundary of f_{max} is a null homologous multicurve made up of horizontal arcs of m_1 to the left of m_2 and horizontal arcs of m_2 to the right of m_1.

Similarly, the subsurface of S, f_{min}, on which f takes on its minimal value is disjoint from f_{max} and is on the left of any arc of m_1 on its boundary and to the right of any arc of m_2 on its boundary.

2.3. **Minimising Overlap.** A difficulty is that vertices of $\mathcal{HC}(S, \alpha)$ are only defined up to isotopy, whereas some of the quantities, such as overlap, used to describe distance also depend on the representative of the homotopy classes. For this reason it is necessary to work with representatives of the free homotopy class that minimise the overlap.

Two multicurves m_1 and m_2 will be said to be in *minimal position* if

- m_1 and m_2 are in general position
- the number of times m_1 intersects m_2 is equal to $i(m_1, m_2)$, and
- whenever $m_2 - m_1$ contains homotopic curves, these homotopic curves are positioned in such a way that f is minimised. An example is illustrated in figure 4.

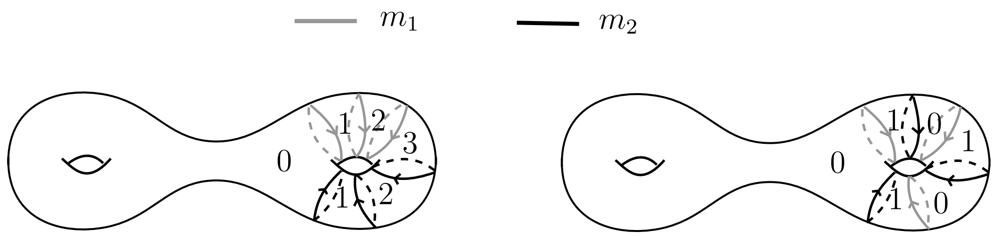

FIGURE 4. On the left, m_1 and m_2 are not in minimal position, because the overlap could be made smaller, as shown on the right.

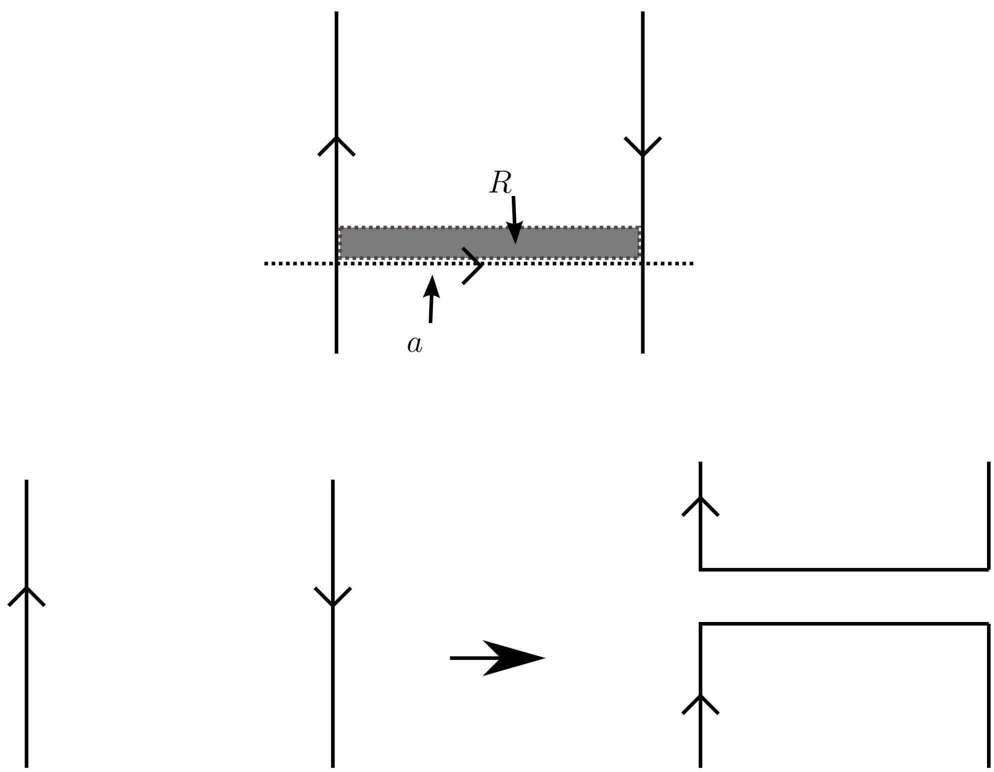

FIGURE 5. Surgering a multicurve along a horizontal arc.

3. A path constructing algorithm

In this section an algorithm for constructing a simple path $m_1, \gamma_1, \gamma_2, ...m_2$ of length $\delta(m_1, m_2)$ will be constructed.

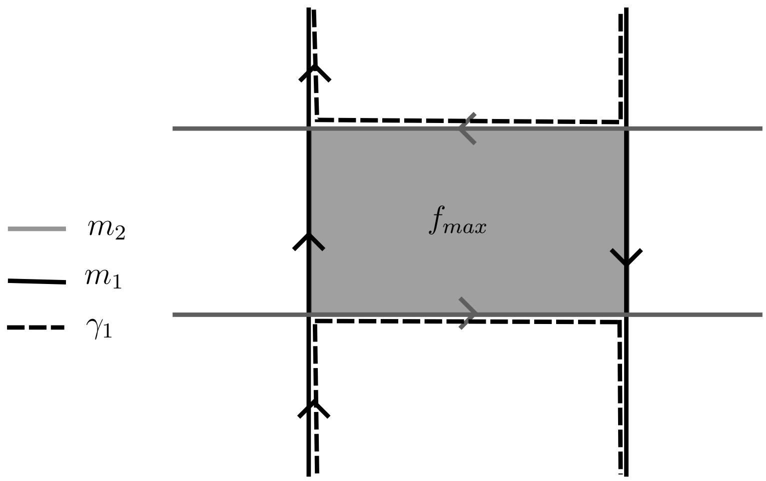

FIGURE 6. How to construct γ_1

A basic surgery construction. Let R be an oriented embedded rectangle in S whose interior is contained in $S \smallsetminus (a \cup b)$. Suppose that one side of R lies along the arc a_i, the opposite side is homotopic to a_i with opposite orientation, and the two remaining sides are subarcs b_1 and b_2 of b, as shown in figure 3. Since a_i is a horizontal arc, it is possible to choose R such that the orientation of R induces an orientation on the arcs b_1 and b_2 on its boundary opposite to the orientation of b_1 and b_2 as subarcs of b. *Surgering an oriented multicurve b along a horizontal arc a_i* of $a \cap (S \smallsetminus b)$ involves adding ∂R to b as a chain. The arcs b_1 and b_2 on the boundary of R cancel out subarcs of b and are replaced by the arcs a_i and $-a_i$. Since ∂R is null homologous, the resulting multicurve is homologous to b.

Recall that the boundary of f_{max} is oriented in such a way that f_{max} is on its left, and let $a_1, a_2...$ be the arcs of m_2 on ∂f_{max}, $b_1, b_2, ...$ be the arcs of m_1 on ∂f_{max}. Then $\partial f_{max} = \sum_i a_i - \sum_j b_j$ (arcs are chains, and so they can be added and subtracted). Consider the one dimensional cell complex $m_1 \cup m_2$ on S. Subtract the oriented arcs b_i from the oriented subcomplex m_1 and add the oriented arcs a_j. This defines γ_1. Subtracting the arcs b_i from m_1 and adding the arcs a_j will be called *performing the surgery* or *surgeries* corresponding to f_{max}, depending on the number of connected components of f_{max}. Up to free homotopy on the boundary, f_{max} can be thought of as "that piece of S that is bounded by m_1 and γ_1".

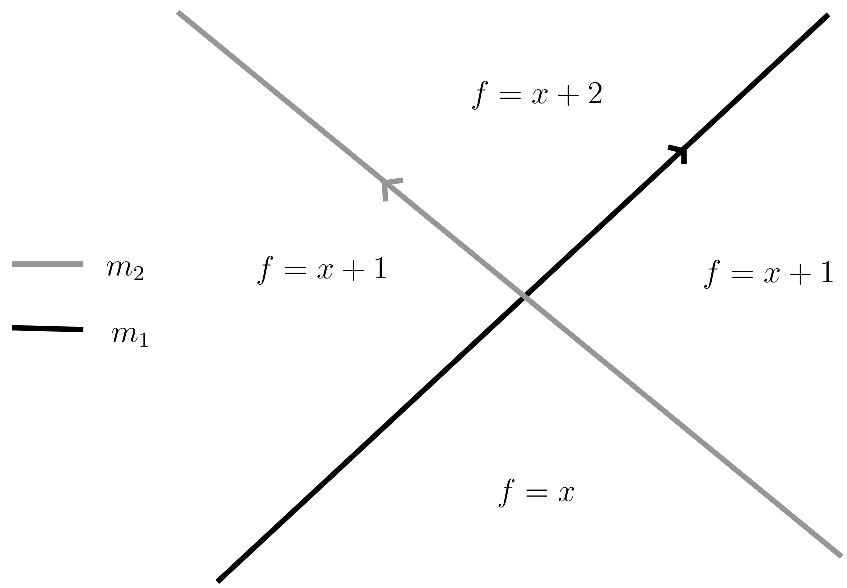

FIGURE 7. A point of intersection forces the overlap to have maximum at least two.

∂f_{max} is disjoint from m_1 and each connected component of f_{max} intersects an annular neighbourhood of m_1 on the right side of m_1 (i.e. every component of f_{max} is "on the same side" of m_1). Therefore $i(\gamma_1, m_1) = 0$.

γ_1 might contain trivial curves that bound disks, and might not be in minimal position with m_2. This point is ignored at the moment, and only once all the multicurves γ_i are constructed are the trivial curves discarded. It follows from the arguments in lemma 6 that despite trivial curves and nonminimal position, none of surgery steps are trivial, i.e. γ_{i+1} is never isotopic to γ_i.

The multicurve γ_2 is constructed in the same way as γ_1 only with the multicurve m_1 replaced by γ_1. It is not difficult to see that the overlap, f_1, of γ_1 and m_2 is one less than the overlap of m_1 and m_2. Cutting out the arcs b_i make it possible to connect the subsurface of S, f_{1min}, on which f_1 takes on its minimum, to f_{1max} (defined similarly), by an arc that crosses $m_2 - \gamma_1$ from right to left once less than any arc connecting f_{min} with f_{max}. In other words, $\delta(\gamma_1, m_2) = \delta(m_1, m_2) - 1$.

This process ends with the multicurve γ_j when $\delta(\gamma_j, m_2) = 1$. This can only be happen if γ_j and m_2 don't intersect, because as shown in figure 7, an intersection forces the maximum of f_j to be at least two.

If $\delta(\gamma_j, m_2) = 1$, then the subsurface of S on which $f_j = 1$ is the subsurface bounded by $m_2 - \gamma_j$.

This completes the construction of the promised algorithm. A simple path constructed in this way will be called a *topdown path*, and the algorithm itself will be referred to as the *path construction algorithm*.

The choice to use f_{max} instead of f_{min} was arbitrary, but it is not possible to simultaneously reduce the intersection number further at each step by requiring that the subsurface of S bounded by γ_1 and m_1 be $f_{max} \cup f_{min}$ because f_{min} is to the left of m_1 and f_{max} is to the right of m_1, so this would not give a simple path.

3.1. **Path Construction in the Cyclic Cycle Complex.** The *Cyclic Cycle Complex* $\mathcal{CC}(S)$ from [8] is the simplicial complex whose vertices are the isotopy classes of oriented, reduced multicurves, where a multicurve m is said to be *reduced* if it does not contain a submulticurve that bounds a complementary region of m in S (using either orientation of the region). A set of $k+1$ vertices spans a simplex in $\mathcal{CC}(S)$ if these vertices are represented by disjoint multicurves m_0, m_1, $m_2 ... m_k$ that cut S into $k+1$ embedded subsurfaces E_0, E_1,...E_k such that the oriented boundary of E_i is $m_{i+1} - m_i$. In particular, all edges are by definition simple.

It follows that each connected component of $\mathcal{CC}(S)$ represents multicurves in a fixed homology class. Every connected component of $\mathcal{CC}(S)$ can therefore be embedded in a $\mathcal{HC}(S, \alpha)$ for appropriate α.

The path construction algorithm can be easily modified to construct paths in $\mathcal{CC}(S)$. This involves removing all null homologous submulticurves that are forbidden by the definition. It is therefore necessary to check that this can be done without violating the condition that a path in $\mathcal{CC}(S)$ has to satisfy. In particular, if a null homologous submulticurve n is to be removed from γ_{i+1}, it is necessary to check that there is a subsurface N of S (with either orientation) with $\partial N = n$ and such that N is disjoint from the subsurface of S bounded by $\gamma_{i+1} - \gamma_i$.

Suppose γ_{i+1} contains a forbidden null homologous submulticurve b_1 i.e. $\pm b_1$ bounds a subsurface B of S, where $\gamma_{i+1} - \gamma_i$ is contained in $S \setminus B$. Whenever this happens, by construction f_{imax} and $\gamma_{i+1} - \gamma_i$ have to be to the left of b_1. Also, whenever γ_i does not contain a forbidden null homologous submulticurve, it follows from the arguments given in the proof of theorem 9 that γ_{i+1} can't contain a second null homologous submulticurve b_2 that lies between b_1 and the other curves in $\gamma_{i+1} - \gamma_i$.

b_1 can therefore be capped off from the right by a subsurface disjoint from f_{imax}. That this path is a geodesic follows from theorem 7.

More General Position. When proving theorems about simple paths, it is convenient to work with multicurves in S that are not in general position. Representatives of the homotopy classes can be chosen to make the intuitive picture of f_{max} clearer. Suppose m_1 and m_2 are in minimal position. Choose the representatives γ_{i+1} and γ_i of $[\gamma_{i+1}]$ and $[\gamma_i]$ such that $\gamma_{i+1} - \gamma_i = \partial f_{imax}$. The boundary of f_{imax} is an embedded subcomplex of the one dimensional cell complex $m_2 \cup m_1$ for every i, and has zero intersection number with γ_i and γ_{i+1}. γ_{i+1} is obtained from γ_i by subtracting the arcs of $\gamma_i \cap (S \setminus m_2)$ on ∂f_{imax} and adding the arcs of $m_2 \cap (S \setminus \gamma_i)$ on ∂f_{imax}. Also, no arc of $m_2 \cap m_1$ will be on the boundary of f_{imax} for more than one i, so each arc can only be added or subtracted at most once. Each of the multicurves γ_i is therefore an oriented subcomplex of $m_1 \cup m_2$. From figure 7, it is easy to verify that f_{imax} can not meet itself at a vertex, because if four components of $S \setminus (\gamma_i \cup \gamma_k)$ come together at a point and f is equal on two of them, it must be larger on a third component and smaller on the fourth. Therefore, if γ_i doesn't meet or cross over itself at a vertex, neither will γ_{i+1}. The γ_i chosen in this way are therefore also embedded. The main advantage of doing this is that the overlap functions $f, f_1, f_2 ...$ are related in an obvious way. In section 6 it will be shown that there does not always exist tight geodesic paths, so geodesic paths can't in general be constructed by adding and subtracting arcs within the cell complex $m_1 \cup m_2$.

A nice property of the path construction algorithm is that it constructs the same path in reverse.

Theorem 3. *If m_1 and m_2 had been interchanged in the path construction algorithm, the same unoriented path would have been obtained.*

Proof. Suppose the representatives $m_1, \gamma_1, ...\gamma_j, m_2$ of the free homotopy classes $[m_1], [\gamma_1], ...[\gamma_j], [m_2]$ are chosen as outlined in the previous paragraph. In particular, each of the γ_i are oriented subcomplexes of the cell complex $m_1 \cup m_2$ such that $\gamma_{i+1} - \gamma_i$ is the boundary of the subsurface of f on which f is no less than its maximum value minus i. Let h be the overlap of $m_1 - m_2$. It is easy to check that h has its maximum where the overlap f of $m_2 - m_1$ has its minimum, and vice versa. By definition, γ_j is the multicurve chosen such that $m_2 - \gamma_j$ bounds the subsurface of S given by $S \setminus f_{min}$. In other words, $\gamma_j - m_2$ is the boundary of f_{min} or h_{max}, i.e. γ_j satisfies the definition of the first multicurve in the path $m_2, ...m_1$. Similarly for γ_{j-1}, γ_{j-2}, etc.

□

4. The Overlap Function and the Pre-image Function

Let H be an oriented, immersed surface H in $S \times I$ with $\pi(\partial H) = m_2 - m_1$. In this section, theorem 7 is proven by relating the overlap of ∂H to the pre-image function g_H.

The *pre-image function* $g_H : S \times 0 \setminus \pi(\partial H) \to \mathbb{Z}$ is defined as follows: Suppose $P := S \times I$ and ∂P is decomposed into the union of two subsurfaces A and B, where A is a neighbourhood in ∂P of ∂H and B is the closure of $P \setminus A$. Algebraic intersection number provides a map $H_2(P, A) \times H_1(P, B) \mapsto \mathbb{Z}$. For x in $(S \times 0) \cap B$, $g_H(x)$ is equal to the algebraic intersection number of $[H]$ with the class in $H_1(P, B)$ represented by the arc $\{x\} \times I$. Since this definition works for any choice of A, $g_H(x)$ is defined for all x in $(S \times 0) \setminus \pi(\partial H)$.

Lemma 4. *Given any two oriented, immersed surfaces H_1 and H_2 with $\partial H_1 = \partial H_2 = m_2 - m_1$, there is a constant integer c such that for all $s \in (S \times 0) \setminus (m_2 - m_1)$, $g_{H_1} = g_{H_2} + c$.*

Proof. Clearly, g_{H_1} and g_{H_2} both increase by one when crossing over an arc of $m_2 - m_1$ from right to left. This lemma is proven by showing that g_{H_1} and g_{H_2} can't change anywhere else. Suppose $\{y\} \times I$ is homologous to $\{x\} \times I$ relative to B, using the same notation as in the definition of the pre-image function. In other words, y and x are in the same component of $(S \times 0) \setminus (m_2 - m_1)$. Then $g_{H_1}(x) = g_{H_2}(y)$, i.e. g_{H_1} is constant on each component of $S \setminus (m_2 - m_1)$. The same argument applies to g_{H_2}, from which the lemma follows.

□

Two surfaces F_1 and F_2 in $S \times I$ with boundary $m_2 - m_1$ are defined to be homotopic if they are homotopic as surfaces with boundaries contained in $m_1 \times I \cup m_2 \times I$. Suppose $m_2 - m_1$ is a multicurve that bounds a subsurface of S and H is any orientable surface in $S \times I$ with boundary $m_2 - m_1$. It follows from the previous lemma that g_H has to be constant on any component of $S \times 0 \setminus (m_2 - m_1)$. From this it follows that if H has smallest possible genus, it has to be homotopic to a subsurface of $S \times 0$.

Corollary 5 (Corollary of lemma 4)**.** *Suppose that $m_2 - m_1$ is a multicurve, where m_1 and m_2 are homologous multicurves in $S \times 0$. Then any orientable surface in $S \times I$ with smallest possible genus whose boundary is homotopic to $m_2 - m_1$ has to be homotopic to a subsurface of $S \times 0$.*

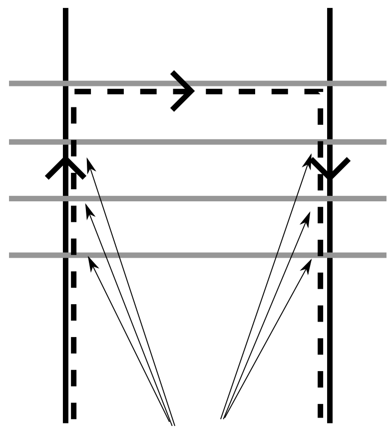

FIGURE 8. Points of intersection that could be removed by a homotopy.

Lemma 6. *Given m_1 and m_2, let γ be the topdown path connecting m_1 and m_2, and let T_γ be its trace. Then the overlap of m_1 and m_2 is equal to g_{T_γ}.*

Proof. The reason this is not immediately clear is that the multicurve γ_{i+1}, obtained from γ_i and m_2 by performing the surgery corresponding to f_{imax}, might contain curves that bound disks or points of intersection with m_2 that can be removed by a homotopy. The overlap f_{i+1} depends on the representative of the free homotopy class $[\gamma_{i+1}]$. In order to define f_{i+1}, it was assumed that the multicurves γ_{i+1} and m_2 were in minimal position.

Whenever a and b are multicurves that are not in general position, i.e. a and b coincide along some subarc or point, this subarc or point will be counted as a (single) point of intersection if b crosses from one side of a to the other. If b does not cross over a this is not counted as an intersection.

Let R_{a_i} be the rectangle in S consisting of the closure of the union of rectangles in $S \setminus (m_1 \cup m_2)$, each of which have two opposite sides made up of arcs of $m_2 \cap (S \setminus m_1)$ in the homotopy class a_i, where each of the $a_1 ... a_n$ are homotopy classes of arcs with representatives on ∂f_{max}. Let $R := R_{a_1} \cup R_{a_2} \cup ... R_{a_n}$, and $\gamma_1^{'}$ be the multicurve homotopic to γ_1 constructed such that $\gamma_1^{'}$ coincides with m_1 outside of R and is a representative of the homotopy class with the smallest possible number of points of intersection with m_2, according to the definition in the previous paragraph. $\gamma_1^{'} - m_1$ therefore bounds the subsurface $f_{max} \cup R$ of S.

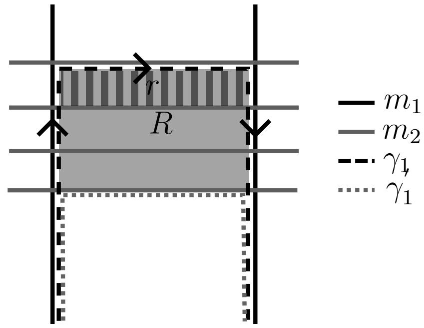

FIGURE 9. γ_1, $\gamma_1^{'}$, R and r.

Let f_1 be the overlap of γ_1 and m_2, and let $f_1^{'}$ be the overlap of $\gamma_1^{'}$ and m_2. γ_1 and $-\gamma_1^{'}$ bound the subsurface R of S. For a point $s \in S$,

$$f_1(s) = \begin{cases} f_1^{'}(s) + 1 & \text{if } s \in \text{the interior of } R, \\ f_1^{'}(s) & \text{otherwise.} \end{cases} \tag{1}$$

In other words, the homotopy that takes γ_1 to $\gamma_1^{'}$ reduces the overlap by one on the subsurface R and enlarges the subsurface of S bounded by $\gamma_1 - m_1$ to obtain the subsurface of S bounded by $\gamma_1^{'} - m_1$, $f_{max} \cup R$.

Any components of $S \setminus (m_1 \cup m_2)$ with one edge along f_{max} are contained in f_{1max}, and since these components aren't all contained in R, it follows that f_1 has the same maximum as $f_1^{'}$. f_{1max} is the union of $f_{1max}^{'}$ with a union r of rectangles of $S \setminus (m_1 \cup m_2)$ in R, as shown in figure 9. The surgeries of γ_1 corresponding to rectangles in r reduce the number of points of intersection with m_2. Homotoping γ_1 to $\gamma_1^{'}$ has the same effect as performing the surgery corresponding to each rectangle in R and discarding contractible curves. If r is not the whole of R, when passing from γ_2 to γ_3, surgeries corresponding to further rectangles in R are performed. This is continued until for large enough i, f_{imax} contains all of R and γ_{i+1} has no points of intersection with m_2 on ∂R. If γ_1 is used in place of $\gamma_1^{'}$ to construct γ_2, the same multicurve will therefore be obtained up to homotopy, despite the fact that γ_1 and m_2 might not be in minimal position. The same argument applies for all γ_i in place of γ_1, from which the lemma follows.

□

It is now possible to give a proof of theorem 7.

Theorem 7. *The shortest simple paths connecting m_1 and m_2 have length $\delta(m_1, m_2)$.*

Proof. Suppose γ is a simple path connecting m_1 and m_2 of length less than $\delta(m_1, m_2)$. Let T_γ be the the trace of γ. Then T_γ can be constructed by connecting up $\delta(m_1, m_2) - 1$ or fewer pieces, each of which projects one to one onto a subsurface of $S \times 0$ with the induced subsurface orientation. By construction, g_{T_γ} is everywhere ≥ 0. It follows from lemmas 4 and 6 that the maximum of g_{T_γ} minus the minimum of g_{T_γ} is equal to $\delta(m_1, m_2)$, i.e $\hat{i}(\pi^{-1}(s), H) \geq \delta(m_1, m_2)$ for some s. This is a contradiction.

From the algortihm given in section 3 it is clear that this minimum length path can always be achieved.

$\square$

Corollary 8 (Corollary of Theorem 7 and Proposition 3.1)**.** *Suppose m_1 and m_2 are two vertices in the Cyclic Cycle Complex. The distance between these two vertices in this complex is equal to $\delta(m_1, m_2)$.*

5. Distances and Simple Paths

Theorem 7 determines the length of the shortest simple paths connecting two vertices, however this has not yet been related to the distance between the vertices. There are two minor points that need to be considered at this point. Firstly, if m_1 contains a null homologous submulticurve b, $m_1 \setminus b$ has the same distance in $\mathcal{HC}(S, \alpha)$ from m_2 as m_1, but $\delta(m_1, m_2)$ could be larger. Secondly, if α is not a primitive homology class, for example, if α is homologous to nm_1, $\delta(nm_1, nm_2) = n\delta(m_1, m_2)$, but the distance between nm_1 and nm_2 in $\mathcal{HC}(S, [nm_1])$ is equal to the distance between m_1 and m_2 in $\mathcal{HC}(S, [m_1])$.

A path between two vertices will be called a *quasi-geodesic (segment)* if it is a subpath of a quasi-geodesic. All quasi-geodesics considered here are *uniform quasi-geodesics*, in the sense that, for any two vertices v_1 and v_2 on the quasi-geodesic, the length of the quasi-geodesic segment jointing them is no more than $kd(v_1, v_2)$, where k is a uniformly bounded constant. Therefore, no distinction between quasi-geodesics and quasi-geodesics segments is made.

Note that, since $\mathcal{HC}(S, \alpha)$ is not δ-hyperbolic (in fact, it is not even nonpositively curved), no geodesic stability should be expected. The quasi-geodesics constructed here do not globally stay close to geodesics, although they can be shown to be piecewise geodesic. Despite this, families of geodesics connecting two vertices in $\mathcal{HC}(S, \alpha)$ can be easily described and constructed due to a high level of rigidity, however this is the subject of a future paper.

Theorem 9. *If m_1 and m_2 do not contain null homologous submulticurves or homotopic curves, the path connecting m_1 and m_2 constructed by the path construction algorithm is a quasi-geodesic, where the constant in the definition of quasi-geodesic is less than $-3\chi(S)$.*

Proof. In the path construction algorithm, the maximum of the overlap was increased by one at each step. This decrease is due to discarding the null homologous submulticurve ∂f_{max} that was created by the surgeries. This theorem is proven by obtaining a bound on the decrease in the maximum of the overlap at each step, by bounding the number of null homologous submulticurves that can created (and therefore potentially discarded) at each step.

It is implicit in the proof of proposition 3 that a path in $\mathcal{HC}(S,\alpha)$ between any two vertices can be constructed by surgering along horizontal arcs and adding/discarding null homologous submulticurves. If a multicurve m does not contain homotopic submulticurves, it follows from the topological invariance of the Euler characteristic that there exists a bound of $-3\chi(S)$ on the number of pairwise disjoint homotopy classes (relative to m) of horizontal arcs with endpoints on m.

Firstly, a proof of the theorem is given in the case that there is a geodesic path $m_1, \gamma_1, \gamma_2, ...m_2$ such that none of the γ_i represent multicurves with homotopic curves. This is done by showing that γ_{i+1} can be obtained from γ_i by surgering along no more than $-3\chi(S)$ horizontal arcs. Since each surgery can increase the number of curves in the multicurve, and hence the number of null homologous submulticurves, by no more than one, the theorem then follows.

Let I be an oriented arc in S that intersects γ_i for some i. There are a certain number of homotopy classes of arcs of $\gamma_i \cap (S \setminus I)$ relative to I. The orientations on I and γ_i makes it possible to define an ordering of the starting points of the arcs of $\gamma_i \cap (S \setminus I)$ along I. Let h be a homotopy of γ_i that changes this ordering without moving any arcs over ∂I. Since γ_i does not contain homotopic curves, h has to introduce self intersections of γ_i. Similarly, if γ_{i+1} also intersects I, then any homotopy of γ_i and/or γ_{i+1} that changes the ordering of the starting points of $\gamma_i \cup \gamma_{i+1}$ along I without moving any arcs over ∂I has to either create (nonessential) points of intersection or move one curve past another curve in the same free homotopy class.

Let $\gamma_{i+1}^{'}$ be the multicurve obtained from γ_i by surgering along the horizontal arcs $a_1, a_2...a_k...a_n$. γ_{i+1} is obtained by discarding null homologous submulticurves from $\gamma_{i+1}^{'}$. Since the path is not assumed to be simple, arbitrarily many null homologous submulticurves might be discarded from $\gamma_{i+1}^{'}$. Assume that at least one of the a_i is of the form

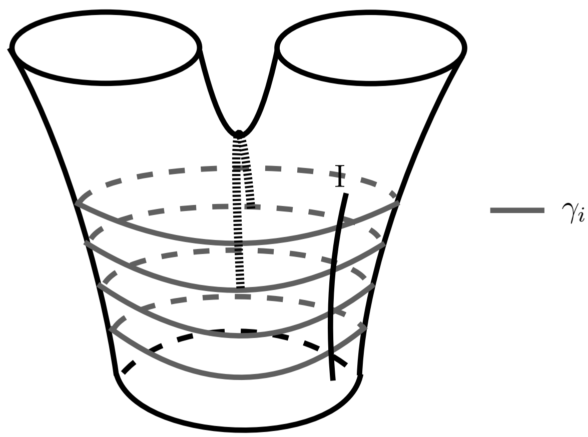

Figure 10. If λ_i could contain homotopic curves, the points of intersection of λ_i with the horizontal arc shown in the figure can be removed by a homotopy that changes the ordering of the points of intersection of λ_i with the interval I, without creating points of self-intersection of λ_i.

$v_1 \circ a_k \circ v_2$ for vertical arcs v_1 and v_2. A contradiction will be obtained, from which it follows that the number of surgeries needed to obtain γ_{i+1} from γ_i is uniformly bounded. Without loss of generality it can also be assumed that none of the surgeries is trivial, i.e. for all i, $S \setminus (\gamma_i \cup a_i)$ is not allowed to contain any bigons. For example, γ_i is not surgered along any two arcs in the same homotopy class. In h there are one or two curves that were created by surgering along a horizontal arc of the form $v_1 \circ a_k \circ v_2$. If γ_{i+1} does not contain at least one of these curves, there was no need to attach the handle corresponding to $v_1 \circ a_k \circ v_2$ at all.

Call a curve in γ_{i+1} *new* if it was created by one of the surgeries in which γ_{i+1} is obtained from γ_i. Either

(1) all new curves in γ_{i+1} are homotopic to other curves in γ_{i+1} i.e. γ_{i+1} contains homotopic curves,
(2) all new curves are homotopic to curves in γ_i, i.e. γ_{i+1} is a submulticurve of γ_i, or
(3) neither 1 nor 2.

Let I be a compact arc in S chosen to pass through an arc in the homotopy class v_1 or v_2. In this third case, if the order of the arcs along I is altered to remove the points of intersection with γ_i of the attached handle corresponding to $v_1 \circ a_k \circ v_2$, it has to induce points

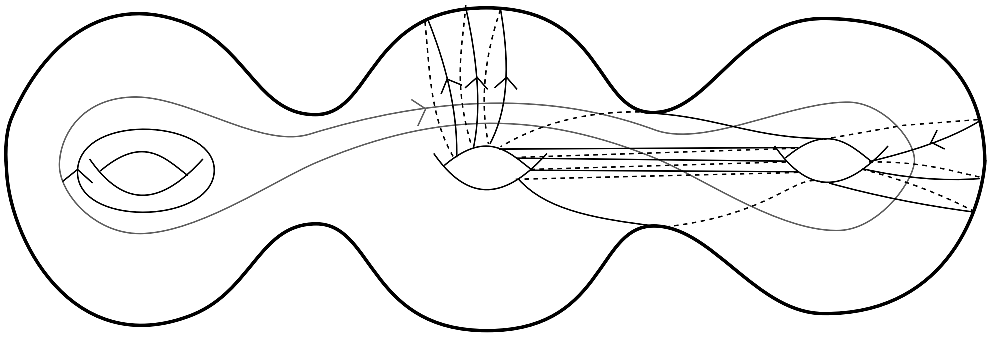

FIGURE 11. A multicurve m homologous to a simple curve (drawn in grey). The multicurve m contains homotopic curves and no null homologous submulticurves.

of intersection elsewhere. In other words, $i(\gamma_{i+1}, \gamma_i) \neq 0$, which is not possible by definition. Since γ_{i+1} is not a submulticurve of γ_i, and by assumption does not contain homotopic curves, the promised contradiction is obtained for a path that doesn't pass through vertices that represent multicurves containing freely homotopic curves.

As shown in figure 11, it is not always possible to get rid of all these homotopic curves by working with multicurves that don't contain null homologous submulticurves.

If the geodesic path $m_1, \gamma_1, \gamma_2, ...m_2$ passes through vertices representing multicurves with homotopic curves, since the γ_i do not contain null homologous multicurves, and m_2 does not contain homotopic or null homologous multicurves, $m_2 - \gamma_i$ does not contain homotopic multicurves that separate f_{imax} from f_{imin}. In other words, there exists an arc connecting a component of f_{imax} to a component of f_{imin} that avoids all homotopic curves of γ_i. Therefore, although arbitrarily many surgeries may be performed on γ_i to obtain γ_{i+1}, this arc constitutes a "bottleneck", to which the same arguments as in the previous case (no homotopic curves) apply.

□

6. Quasi-flats and Distance Bounds

This section gives a few simple examples to illustrate key geometric properties of $\mathcal{HC}(S, \alpha)$.

Distances in $\mathcal{HC}(S, \alpha)$ were shown to be related the homological distance. The next question is, how does this relate to intersection number? At each step of the path construction algorithm, the intersection

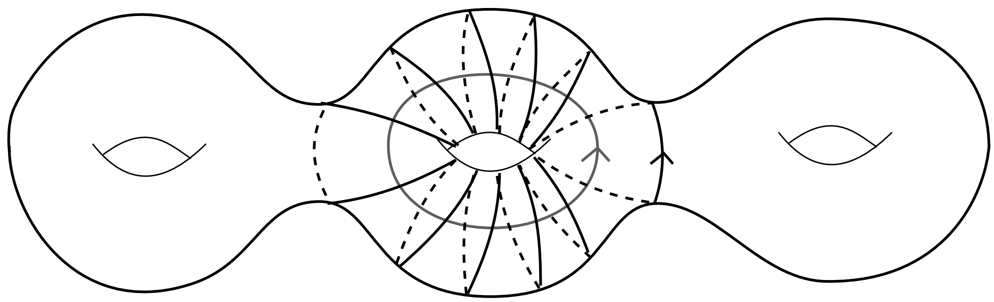

FIGURE 12. Example demonstrating that the best possible upper bound on the distance between c_1 and c_2 in $\mathcal{HC}(S, \alpha)$ is given by $\frac{i(c_2,c_1)}{2} + 1$.

number with m_2 is decreased. Recall that the arcs of $m_2 \cap (S \setminus m_1)$ on ∂f_{max} were denoted $a_1 ... a_n$. Let k_{a_i} be the number of arcs of $m_2 \cap (S \setminus m_1)$ in the same homotopy class as a_i for $1 \leq i \leq n$. Then the intersection number of γ_1 with m_2 is at least $2 \sum_i k_{a_i}$ less than the intersection number of m_1 with m_2.

It is well known that the distance between two curves c_1 and c_2 in the curve complex is either 1 if $i(c_1, c_2) = 0$ or is bounded from above by $\log_2(i(c_1, c_2)) + 1$. The next example shows that the distance between two curves in $\mathcal{HC}(S, \alpha)$ can be as much as $\frac{i(c_1,c_2)}{2} + 1$.

Example 10 (Dehn twisting around bounding pairs)**.** *Let c_1 and c_2 be the curves shown in figure 12. c_2 is obtained by Dehn twisting c_1 n times around a bounding pair, where $n = 5$ in figure 12. A simple calculation shows that $\delta(c_1, c_2)$ is equal to $\frac{i(c_1,c_2)}{2} + 1$. In this case, it is also clear that $\frac{i(c_1,c_2)}{2} + 1$ is the distance between c_2 and c_1 in $\mathcal{HC}(S, \alpha)$. To see why, note that any multicurve in α has to have nonzero algebraic intersection number with each of the curves in the bounding pair. Also, it is not possible to Dehn twist more than once around the bounding pair when passing from γ_i to γ_{i+1}. From this it follows that $i(\gamma_{i+1}, c_2) \geq i(\gamma_i, c_2) - 2$, i.e. a shorter path than the path obtained from the path construction algorithm can not exist.*

Unlike the curve complex, which is known to be δ-hyperbolic ([13] and [2]), this observation can be used to provide an example to show that $\mathcal{HC}(S, \alpha)$ is not δ-hyperbolic.

Theorem 12. *$\mathcal{HC}(S, \alpha)$ is not δ-hyperbolic for $g > 3$.*

Proof. For $g > 3$ there exist two pairs of bounding pairs (t_1, t_2) and (t_3, t_4); each of the t_i representing distinct isotopy classes. Suppose

Example 11 (Dehn twisting around pairs of bounding pairs).

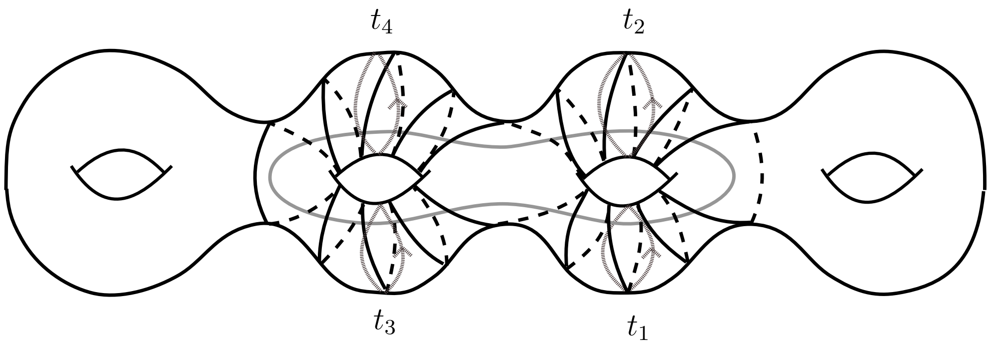

v_1 is a multicurve with nonzero algebraic intersection number with each of t_1, t_2, t_3 and t_4, as in example 11. Let v_2 be the multicurve v_1 Dehn twisted around (t_1, t_2) n times, and v_3 be the multicurve v_1 Dehn twisted around (t_3, t_4) n times. v_1, v_2 and v_3 represent the vertices of a geodesic triangle in $\mathcal{HC}(S, \alpha)$. Since the distance between two vertices on the boundary of the triangle is equal to the number of Dehn twists around the bounding pairs (t_1, t_2) and (t_3, t_4) necessary to get from one vertex to the other, for n even, the midpoints of the sides of the geodesic triangle are each a distance $\frac{n}{2}$ from the other two sides of the triangle. For any fixed δ, n can therefore be chosen large enough so that this triangle is not δ-thin. In this example, the triangle is contained in a so-called quasi-flat. □

Example 11 also shows that, unlike in the curve complex, there does not always exist a tight geodesic connecting any two vertices. A geodesic $c_1, \gamma_1, \gamma_2, ..., c_2$ has to be constructed such that for each i, γ_{i+1} is obtained from γ_i by performing four Dehn twists. It is not hard to check that this is only possible if γ_1 is obtained from c_1 by performing a surgery that cuts c_1 into two curves; one that intersects t_1 and t_2, and another one that intersects t_3 and t_4. All curves contained in the one dimensional cell complex $c_1 \cup c_2$ are either null homologous, c_1, c_2, t_1, t_2, t_3, t_4 or they intersect all of t_1, t_2, t_3 and t_4. It follows that a geodesic connecting c_1 and c_2 can't be tight.

References

[1] Bestvina, M. Bux, K. and Margalit, M. (2007) 'The dimension of the Torelli group' arXiv:0709.0287.

[2] Bowditch, B. (2006), 'Intersection numbers and hyperbolicity of the curve complex' *J. reine angew. Math* **598**, 105-129

[3] Bowditch, B. (2008), 'Tight geodesics in the curve complex' *Invent. math.* **171**, 281-300.

[4] Brock, J., Canary, R. and Minsky, Y. (2004), 'The classification of Kleinian groups, II: The ending lamination conjecture' arXiv:math/0412006v1.

[5] Farb, B. and Margalit, D. (2009) *A Primer on Mapping Class Groups version 3.1* available at www.math.uchicago.edu/ margalit/mcg/mcgv31.pdf

[6] Farb, B. and Ivanov, N. (2003) 'The Torelli geometry and its applications' Research Announcement.

[7] Harvey, W. (1981), 'Boundary structure of the modular group' Riemann Surfaces and Related Topics: Proceedings of the 1978 Stony Brook Conference (I. Kra and B. Maskit, eds), Ann. of Math. Stud. 97, Princeton

[8] Hatcher, A. (2008), 'The cyclic cycle complex of a surface', available at http://www.math.cornell.edu/ hatcher/Papers/cycles

[9] Hempel, J. (2001), '3-manifolds as viewed from the curve complex' Topology **40**(3), pp. 631-657.

[10] Irmer, I. (2010),'The curve graph and surface construction in $S \times \mathbb{R}$', PhD Thesis.

[11] Leasure, J. (2002), 'Geodesics in the complex of curves of a surface' PhD Thesis. Available at http://hdl.handle.net/2152/1700

[12] Lickorish, W. (1964), 'A finite set of generators for the homeotopy type of a 2-manifold' Math. Proc. Cambridge Phil. soc. **60**.

[13] Masur, H., Minsky, Y.(1999), 'Geometry of the complex of curves I: Hyperbolicity' *Invent. Math.***138**, 103-149.

[14] Masur, H., Minsky, Y. (2000), 'Geometry of the complex of curves II: Hierarchical Structure' *Geometric and Functional Analysis***10(4)**

[15] Shackelton, J. 'Tightness and computing distances in the curve complex' arXiv:math/0412078v3 14 Apr 2005

E-mail address: irmer@cs.uni-bonn.de