

An acylindricity theorem for the mapping class group

Kenneth J. Shackleton

ABSTRACT. We study the usual action of the mapping class group of a surface on the 1-skeleton of Harvey’s curve complex from a computational perspective, finding that the number of mapping classes moving a “long” geodesic path a “small” distance is explicitly bounded in terms of intersection numbers and the surface topological type.

1. Introduction

Let Σ be a compact, connected and orientable surface of genus $g(\Sigma)$ and $\#\partial\Sigma$ boundary components. In (7), Harvey associates to Σ a simplicial complex $\mathcal{C}(\Sigma)$ called the *curve complex*. As well as encoding some of the asymptotic geometry of the Teichmüller metric, the curve complex plays a central role in the celebrated proof of Minsky and his collaborators of Thurston’s ending lamination conjecture (12; 3). A key step in their strategy is a theorem of Masur-Minsky’s (10), that every curve graph is hyperbolic in the sense of Gromov.

The curve complex is constructed as follows. We say that a simple loop on Σ is *trivial* if it bounds a disc and *peripheral* if it bounds an annulus containing one component of $\partial\Sigma$. Let $X(\Sigma)$ be the set of all free homotopy classes of non-trivial and non-peripheral simple loops on Σ . The elements of X are referred to as *curves*. We take the vertex set of \mathcal{C} to be X and deem a family of pairwise distinct curves to span a simplex if and only if any two of its curves admit disjoint representatives. Any maximal set of pairwise distinct and pairwise disjoint curves has cardinality $\xi(\Sigma) := \max\{0, 3g(\Sigma) + \#\partial\Sigma - 3\}$.

With the exception of only seven cases, namely either Σ is a sphere and $\#\partial\Sigma \leq 4$ or Σ is a torus and $\#\partial\Sigma \leq 1$, the set $X(\Sigma)$ is non-empty and the curve complex is connected. For these non-exceptional cases, it can be verified that the simplicial dimension of $\mathcal{C}(\Sigma)$ is equal to $\xi(\Sigma) - 1$. We see that Σ is in fact non-exceptional if and only if $\xi(\Sigma) \geq 2$.

Mathematics Subject Classification. 57M50 (primary); 30F60 (secondary).

Key words and phrases. surface; curve complex; mapping class group.

The author wishes to thank the Centre Interfacultaire Bernoulli EPFL for its hospitality and to thank Brian Bowditch for many interesting conversations. The author also wishes to thank the referee for reading this paper and for making many useful suggestions.

When this is the case, the curve complex can be endowed with the canonical path-metric by first declaring each edge to have length equal to 1 and then by taking euclidean simplices. All that is important in this paper is the 1-skeleton of the curve complex, denoted \mathcal{G} , which in its own canonical path-metric is quasi-isometric to \mathcal{C} via the natural inclusion. We refer to the simplicial graph \mathcal{G} as the *curve graph*. All distances will be taken in this graph, whose own canonical metric we denote by d . We remark the curve graph is nowhere locally finite and, when $\xi(\Sigma) \geq 3$, between two vertices of distance at least 2 there often exist infinitely many geodesic paths.

The *mapping class group* of Σ , denoted $\text{Map}(\Sigma)$, we define as the group of all self-homeomorphisms of Σ modulo the normal subgroup of those homeomorphisms homotopic to the identity. This group has a natural cocompact action on the curve graph and this has been exploited by numerous authors in numerous ways; see (6) or (8) by way of example.

The purpose of this paper is to prove a new theorem regarding the nature of this action. There have been at least two studies along these lines already, beginning with the work of Bestvina-Fujiwara (1) who, prompted by an argument of Luo's sketched immediately after Proposition 4.6 of (10) and itself a moderate simplification of an argument of Kobayashi's (9), establish a certain "weak proper discontinuity (WPD)" property. This is an important result, for Bestvina-Fujiwara use this to prove that the dimension of the second bounded cohomology of the mapping class group, and any one of its non-virtually abelian subgroups, is infinite.

Inspired by their arguments, Bowditch (2) proves the *acylindricity* of this action: the number of mapping classes moving a "long" geodesic a "small" distance is uniformly bounded in terms of the topology of the surface. This is stronger than WPD, where the number of such mapping classes is assumed to be finite and where one end of the geodesic is the image of the other under a pseudo-Anosov mapping class raised to a sufficiently high power. Many of the interesting groups that admit an acylindrical or WPD action on a hyperbolic metric space necessarily contain free groups of rank 2 and thus of arbitrarily high rank; see (4).

At some stage, both proofs from (1) and (2) make essential use of passing from a sequence of curves to a limiting lamination to ultimately derive a contradiction and, as such, all the information found this way would appear to be non-computable. The main result of this paper may be viewed as a computational alternative to the acylindricity theorem of Bowditch.

Theorem 1. *Suppose Σ is non-exceptional. There exists a computable function $F: \mathbb{N}^2 \rightarrow \mathbb{N}$ and a constant C such that the following holds: Let r be any non-negative integer. Then, for any two curves α and β with $d(\alpha, \beta) \geq (C + 2)r + 9$, the number of mapping classes $h \in \text{Map}(\Sigma)$ satisfying $d(\alpha, h\alpha) \leq r$ and $d(\beta, h\beta) \leq r$ is bounded above by $F(\iota(\alpha, \beta), r)$.*

We remark the lower bound on $d(\alpha, \beta)$ can likely be improved. Even so, this is stronger than the WPD property found by Bestvina-Fujiwara, and logically independent of Bowditch's acylindricity theorem. While the

bound on the number of such mapping classes does depend on intersection number, whereas the bound given in Bowditch’s acylindricity theorem does not, in contrast it is computable and our argument is entirely elementary. Of much interest is an argument that overcomes this trade-off, yielding both computability and uniformity simultaneously.

2. Background and definitions

We begin by recalling several definitions, the most important of which is that of a “tight multigeodesic” after Masur-Minsky.

2.1. Curves and multicurves. Given any two curves α and β , their *intersection number* $\iota(\alpha, \beta)$ is defined equal to $\min\{|a \cap b| : a \in \alpha, b \in \beta\}$. A *multicurve* is a non-empty collection of pairwise distinct curves of pairwise zero intersection number, and as such spans a simplex in the curve complex. The intersection number of two multicurves is defined additively.

2.2. Paths and multipaths. We regard a *path* in the curve graph as a non-empty sequence of vertices in which consecutive vertices span an edge. We denote the length of a path p , defined equal to the number of vertices less 1, by $l(p)$. A *geodesic* in the curve graph is then a path whose length is precisely the distance between its ends, or terminal vertices.

A *multipath* (ν_0, \dots, ν_n) is a sequence of multicurves such that $(\gamma_0, \dots, \gamma_n)$ is a path for each curve $\gamma_i \in \nu_i$ over each index $i \in \{0, \dots, n\}$. We refer to each multicurve ν_i as a *vertex* of the multipath, following the language of Masur-Minsky. We denote the length of a multipath p , also defined equal to the number of vertices less 1, by $l(p)$. A *multigeodesic* is a multipath (ν_0, \dots, ν_n) such that $(\gamma_0, \dots, \gamma_n)$ is a geodesic, for each curve $\gamma_i \in \nu_i$ over each index $i \in \{0, \dots, n\}$.

We say a multipath (ν_0, \dots, ν_n) is *k-embedded* for a non-negative integer k if, for any two indices i and j with $|i - j| \geq k$, we have $d(\gamma_i, \gamma_j) \geq k$ for each curve $\gamma_i \in \nu_i$ and each curve $\gamma_j \in \nu_j$. Any k -embedded multipath (ν_0, \dots, ν_n) of length at least k is locally k -geodesic, in the sense $d(\gamma_i, \gamma_j) = |i - j|$ for $|i - j| \leq k$ and any $\gamma_i \in \nu_i$ and $\gamma_j \in \nu_j$. Any multigeodesic is k -embedded for each non-negative integer k .

A multipath (ν_0, \dots, ν_n) is said to be *k-quasi-geodesic* for a non-negative integer k if $n \leq d(\gamma_0, \gamma_n) + k$ for each curve $\gamma_0 \in \nu_0$ and each curve $\gamma_n \in \nu_n$. In particular, a path is k -quasi-geodesic if and only if the same path is $(1, k)$ -quasi-geodesic in the standard sense of (5).

Finally, we say a set of multicurves p *extends* to a multipath q if as sets $p \subseteq q$. For any set of multicurves p and any subset $B \subseteq X(\Sigma)$, we define $p - B$ to be equal to the set $\{\nu \in p : \nu \cap B = \emptyset\}$. Note $p - B$ extends to p .

2.3. Tight multigeodesics. The notion of a tight multigeodesic was introduced by Masur-Minsky (11) to address the lack of local finiteness in the curve graph. Though there often exist infinitely many geodesics connecting a pair of vertices of distance at least 2, whenever $\xi(\Sigma) \geq 3$, Corollary 6.4 of (11) states that the number of tight multigeodesics connecting any given pair of vertices is always finite. A slightly weaker definition was later offered

by Bowditch in (2), where the finiteness result of Masur-Minsky is strengthened. In (13; 14), the author offers computable bounds on the number of tight multigeodesics, in either definition, connecting any pair of vertices.

We work exclusively with Masur-Minsky's definition, recalled as follows. For any two multicurves ν_0 and ν_2 connected by a multigeodesic of length 2, we realise ν_0 and ν_2 by transverse representatives intersecting $\iota(\nu_0, \nu_2)$ times and denote by N an open regular neighbourhood of their union. Attach to N all the discs, and all the one-holed discs containing a component of $\partial\Sigma$, complementary to N . The free homotopy class of the non-peripheral boundary components of the resulting surface is a well-defined multicurve associated to ν_0 and ν_2 , so long as we disregard multiplicity. We denote this multicurve by $\partial(\nu_0, \nu_2)$, and refer to it as the *relative boundary of ν_0 and ν_2* . Note every curve in $\partial(\nu_0, \nu_2)$ is distinct and disjoint from every curve in $\nu_0 \cup \nu_2$. An example of a relative boundary ν_1 is given below in Figure 1.




FIGURE 1. The relative boundary of ν_0 and ν_2 .

We now recall the relevant definition of a tight multigeodesic.

Definition 2 (Masur-Minsky). *A multigeodesic (ν_0, \dots, ν_n) is said to be tight at index j , for $j \in \{1, \dots, n-1\}$, if $\nu_j = \partial(\nu_{j-1}, \nu_{j+1})$. We say that (ν_0, \dots, ν_n) is tight if it is tight at each such index.*

Lemma 4.5 of (11) states the existence of at least one tight multigeodesic connecting any two given vertices of the curve graph, found as follows. Suppose $(\nu_0, \nu_1, \nu_2, \nu_3)$ is a multigeodesic tight at ν_2 . We replace ν_1 with the relative boundary of ν_0 and ν_2 . It is verified in the proof of Lemma 4.5 from (11) that this does not affect tightness at ν_2 , so that $(\partial(\nu_0, \nu_2), \nu_2, \nu_3)$ is a tight multigeodesic. Moreover, we are free to replace ν_0 with a second multicurve ν'_0 without affecting tightness at ν_2 in any way, so long as $(\nu'_0, \nu_1, \nu_2, \nu_3)$ is a multigeodesic. This observation is implicit in the fourth paragraph of the proof of Lemma 4.5 from (11).

We summarise this discussion as follows.

Lemma 3 (Masur-Minsky). *Let (ν_1, ν_2, ν_3) be a tight multigeodesic, and let ν_0 denote any multicurve such that $(\nu_0, \nu_1, \nu_2, \nu_3)$ is a multigeodesic. Then, $(\nu_0, \partial(\nu_0, \nu_2), \nu_2, \nu_3)$ is a tight multigeodesic.*

The tightening procedure is very robust, and can be applied to the vertices of a given multigeodesic in any order to produce a tight multigeodesic with the same ends. As observed in (11) the tight multigeodesic that results from tightening at each non-terminal vertex of a geodesic path seems to depend on the order in which the vertices are tightened.

As we see in this paper, the tightening procedure can be generalised to accept a larger collection of multipaths. While we also have to keep in mind the technical difficulties of working with multicurves, and not just single curves, the main issue is the multipath that results from tightening a 3–embedded multipath at a single vertex might not be 3–embedded. This, however, presents the opportunity to locally shortcut the resulting multipath. Combining the two operations of tightening and shortcutting, we have a finite-time algorithm which returns a tight and 3–embedded multipath sharing the same ending vertices as the input multipath. What is more, “much” of the original multipath will remain intact if it happens to be geodesic and tight over much of its length. The full details of this procedure amount to the proof of Lemma 4 in Section 3.

2.4. Hyperbolicity. We recall two definitions from (5). A metric space (X, d) is c -hyperbolic for some non-negative real number c if for any geodesic triangle the c -regular neighbourhood any two sides contains the third side. A subset $B \subseteq X$ is c -quasi-convex if any two points of B are connected by a geodesic path contained in the c -regular neighbourhood of B .

3. Proof of Theorem 1

Let us introduce the notation $B(\alpha, r)$ for the ball of radius $r \in \mathbb{N}$ in the curve graph centred on the vertex $\alpha \in X(\Sigma)$, and $B(\alpha, \beta; r)$ for the union $B(\alpha, r) \cup B(\beta, r)$. Let us also introduce the notation $T(\alpha, \beta)$ for the set of all curves belonging to vertices of tight multigeodesics connecting α to β . For r a non-negative integer, we define $T(\alpha, \beta; r)$ to be equal to the union of all sets $T(\delta, \gamma)$ where $\delta \in B(\alpha, r)$ and $\gamma \in B(\beta, r)$. For a second non-negative integer s , we define $T(\alpha, \beta; r, s)$ to be equal to the set

$$\{\eta \in T(\alpha, \beta; r) : d(\alpha, \eta) \geq s, d(\eta, \beta) \geq s\}.$$

The supporting Lemma 4 offers computable bounds on the number of curves that can lie on a tight multigeodesic, sufficiently far from its ends, that connects two bounded subsets sufficiently far apart in the curve graph. In what follows, let c be any choice of hyperbolicity constant and let C be the constant defined equal to $61 + 1000c$.

Lemma 4. *Suppose Σ is non-exceptional. There exists a computable function $F_1: \mathbb{N}^2 \rightarrow \mathbb{N}$ such that the following holds: For r any non-negative integer and α and β any two curves, $|T(\alpha, \beta; r, Cr)| \leq F_1(\iota(\alpha, \beta), r)$.*

Before proving Lemma 4, let us complete the proof of Theorem 1.

Proof [of Theorem 1] We start by connecting α to β by a tight multigeodesic z . For any mapping class $h \in \text{Map}(\Sigma)$ such that $d(\alpha, h\alpha) \leq r$

and $d(\beta, h\beta) \leq r$, the h -translate of z is a tight multigeodesic whose curves all belong to the set $T(\alpha, \beta; r)$. In particular, our choice of lower bound for $d(\alpha, \beta)$ implies the existence of a pair of curves $\{\delta, \eta\}$ contained in the vertices of z and of distance 3 such that each of their translates by each such mapping class h is contained in $T(\alpha, \beta; r, Cr)$. According to Lemma 4, the number of such translates of either δ or η is at most $F_1(\iota(\alpha, \beta), r)$. Combining this fact with Lemma 7.4 from (2), asserting that the stabiliser in $\text{Map}(\Sigma)$ of a pair of curves at least distance 3 apart in the curve graph is uniformly and explicitly bounded in terms of $\xi(\Sigma)$, we complete a proof of Theorem 1. \square

To prove Lemma 4 we rely on the following theorem, a proof of which is implicit in the proof of Theorem 1.3 from (13) and in (14). The argument found therein proceeds by contradiction, and constructs a path of length at most 2 connecting two curves of distance at least 3 and as such needs only the 3-embedded property of tight multigeodesics to apply.

Theorem 5. *Suppose Σ is non-exceptional. There exists a computable function $F_2 : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that the following holds: Let k denote any non-negative integer. For any k -quasi-geodesic, 3-embedded and tight multipath (ν_0, \dots, ν_n) , both $\iota(\nu_0, \nu_j)$ and $\iota(\nu_j, \nu_n)$ are at most $F_2(\iota(\nu_0, \nu_n), k)$ for each index j .*

Proof [of Lemma 4] For any multigeodesic z beginning in $B(\alpha, r)$ and ending in $B(\beta, r)$, suppose there exists a tight $2r$ -quasi-geodesic multipath q connecting α to β , containing every curve from z at least distance Cr from both α and β , and which is 3-embedded. According to Theorem 5, each vertex of q has intersection number with α and intersection number with β uniformly and explicitly bounded in terms of the quasigeodesic parameter of q , and hence in terms of r , and the intersection number $\iota(\alpha, \beta)$. Each vertex of z at least distance Cr from both α and β thus has similarly bounded intersection number with α and with β . As α and β together fill Σ , this is enough to explicitly bound the cardinality of $T(\alpha, \beta; r, Cr)$.

All that remains is to establish the existence of such a multipath, whose length is at most $d(\alpha, \beta) + 2r$, and we do so by a careful surgery argument. To this end we will introduce two new and non-symmetric operations between multipaths, denoted **S** and **T**.

The operation **S** will describe the shortening of a given multipath where it fails to be 3-embedded. The operation **T** will describe the tightening of a 3-embedded multipath at a single vertex. When the original multipath is both geodesic and tight over much of its length these operations tend to be localised and much of the original multipath remains intact. In particular, this is the case when our multipath is the concatenation of a “short” multi-geodesic, a “long” tight multigeodesic, and then another short multigeodesic each ending on single curves.

The multipath that results from tightening a 3-embedded multipath once might not be 3-embedded and will be in need of further shortening. However, such complications only occur a bounded number of times, for path length strictly decreases upon each **S** and the number of non-tight vertices decreases with each **T**. Thus our tightening procedure will soon stabilise,

indeed in quadratic time, at which point we have found a tight and 3–embedded multipath connecting the same two vertices.

We now give formal definitions. For two multipaths p and p' with common ends we write $p\mathbf{S}p'$ if there exist consecutive vertices (μ_i, \dots, μ_j) of p , with $i \leq j - 2$, and consecutive vertices $(\omega_i, \dots, \omega_k)$ of p' , with $0 \leq k - i \leq \min\{2, j - i - 1\}$, such that:

- $\omega_i, \omega_k \in X(\Sigma)$ are both curves;
- $(\omega_i, \dots, \omega_k)$ is a tight multigeodesic, called the *shortcut*;
- $p \setminus \{\mu_i, \dots, \mu_j\} = p' \setminus \{\omega_i, \dots, \omega_k\}$;
- $\omega_i \in \mu_i$, and
- $\omega_k \in \mu_j$.

Note that the length of p is strictly greater than the length of p' , thus $p \neq p'$, and that p' may fail to be tight at the two vertices of p' either side of the shortcut $(\omega_i, \dots, \omega_k)$ and at both ω_i and ω_k . An instance of the operation \mathbf{S} is depicted below in Figure 2.




FIGURE 2. The multipath p' is formed by shortcutting p , so $p\mathbf{S}p'$.

For a 3–embedded multipath q and a multipath q' with the same ends but distinct from q , we write $q\mathbf{T}q'$ if q' results from tightening q at a single vertex. We note that even if $q\mathbf{T}q'$, while q is by definition 3–embedded we do not presume q' is 3–embedded. If $q\mathbf{T}q'$, then q and q' have the same length and q' is tight at an additional vertex.

In what follows, we define a constant K equal to $50c + 3$. We fix an arbitrary non-negative integer r , and two curves α and β such that $d(\alpha, \beta) \geq 2r + 5$. For a multipath p we denote $p - B(\alpha, \beta; s)$ by $p[s]$ for any non-negative integer s , noting $p[t] \subseteq p[s]$ whenever $s \leq t$.

Lemma 6. *For two multipaths p and p' , if $p[r]$ extends to a multigeodesic and $p\mathbf{S}p'$ then $p[r + K] \subseteq p'$.*

Proof Suppose otherwise, for contradiction. Then, there exists a vertex $\nu \in p[r + K]$ such that $\nu \notin p'$. As $p[r]$ extends to a multigeodesic, the shortcut

does not connect two curves from two distinct vertices of $p[r]$. Moreover ν is not contained in a multigeodesic whose ends both intersect either $B(\alpha, r+1)$ or $B(\beta, r+1)$ as all balls are c -quasi-convex. It follows the shortcut does not connect two curves in $B(\alpha, r+1)$ or two curves in $B(\beta, r+1)$. The shortcut can therefore only connect a curve from $B(\alpha, r+1)$ to a curve from $B(\beta, r+1)$, or connect a curve from a vertex $\nu' \in p[r]$ to a curve in $B(\gamma, r+1)$ where $\gamma \in \{\alpha, \beta\}$ is such that if $\nu' \neq \nu$ then ν separates ν' along p from $B(\gamma, r+1)$. Now the former case is absurd, for the shortcut has length at most 2 while every curve from $B(\alpha, r+1)$ is at least distance $d(\alpha, \beta) - 2(r+1) \geq 2r + 5 - 2r - 2 = 3$ from any curve in $B(\beta, r+1)$. In the latter case, for any curve $\delta \in \nu$ and any curve $\delta' \in \nu'$, the hyperbolicity of the curve graph gives $d(\delta', \gamma) \geq d(\delta, \gamma) - 20c \geq (r + 50c + 4) - 20c = r + 30c + 4 \geq r + 4$. In particular, δ' is at least distance 3 from every curve in $B(\gamma, r+1)$. This is a contradiction. \square

Lemma 7. *For two multipaths p and p' , if $p[r]$ extends to a multigeodesic and $p\mathbf{S}p'$ then $p'[r+3] \subseteq p[r]$.*

Proof Suppose otherwise, for contradiction. Then, there exists a vertex $\nu \in p'[r+3]$ that does not belong to $p[r]$. In particular, ν is contained in the shortcut forming p' from p . As $p[r]$ extends to a multigeodesic, the shortcut does not connect two curves from distinct vertices of $p[r]$. The shortcut therefore either connects two curves in $B(\alpha, \beta; r+1)$, or connects a curve in a vertex of $p[r]$ to a curve in $B(\alpha, \beta; r+1)$. In either case we find $\nu \cap B(\alpha, \beta; r+3)$ is not empty and so $\nu \notin p'[r+3]$. This is a contradiction. \square

Lemma 8. *For a 3-embedded multipath q and a multipath q' , if $q[r]$ is tight and $q\mathbf{T}q'$ then $q[r+1] \subseteq q'$.*

Proof Let ν be a vertex of $q[r+1]$. If ν is not a terminal vertex of $q[r]$ then ν is the relative boundary of two vertices of $q[r]$, in which case ν belongs to q' . If ν is a terminal vertex of $q[r]$ then ν is also a terminal vertex of q and belongs to q' . \square

Lemma 9. *For a 3-embedded multipath q and a multipath q' , if $q[r]$ is tight and $q\mathbf{T}q'$ then $q'[r+1] \subseteq q$.*

Proof Suppose otherwise, for contradiction. Then, there exists a vertex $\nu \in q'[r+1]$ that does not belong to q . Now ν is the relative boundary of two vertices of q and, since $q[r]$ is tight, one of these vertices must have non-empty intersection with $B(\alpha, \beta; r)$. It follows ν has non-empty intersection with $B(\alpha, \beta; r+1)$ and as such does not belong to $q'[r+1]$. This is a contradiction. \square

Given two multipaths p_0 and p_n and a finite word $\mathbf{W} = \mathbf{A}_1 \cdots \mathbf{A}_n$, where $\mathbf{A}_i \in \{\mathbf{S}, \mathbf{T}\}$ for each $i \in \{1, \dots, n\}$, we write $p_0\mathbf{W}p_n$ if there exist multipaths p_1, \dots, p_{n-1} such that $p_0\mathbf{A}_1p_1 \cdots \mathbf{A}_{n-1}p_{n-1}\mathbf{A}_np_n$. As path length and the number of non-tight vertices decrease with each \mathbf{S} and \mathbf{T} respectively, whenever $p\mathbf{W}q$ the length $l(\mathbf{W})$ of \mathbf{W} is certainly bounded above by $l(p)^2$, the square of the length of p . Finally, we say that the word \mathbf{W} is *compatible with p* if there exists a multipath q such that $p\mathbf{W}q$.

Lemma 10. *For any multipath p , any word \mathbf{W} in $\{\mathbf{S}, \mathbf{T}\}$ compatible with p , if $p\mathbf{W}q$ for a multipath q and $p[r]$ extends to a tight multigeodesic then $p[r + Kl(\mathbf{W})] \subseteq q$.*

Proof We argue by induction on word length. If \mathbf{W} is the empty word then $l(\mathbf{W})$ is zero and $q = p$, completing the base case. For the inductive step, we may express \mathbf{W} as a product \mathbf{UA} where \mathbf{U} is a word in $\{\mathbf{S}, \mathbf{T}\}$ of length $l(\mathbf{W}) - 1$ and $\mathbf{A} \in \{\mathbf{S}, \mathbf{T}\}$. As \mathbf{W} is compatible with p there exist multipaths q and q' such that $p\mathbf{U}q\mathbf{A}q'$. The inductive hypothesis yields $p[r + Kl(\mathbf{U})] \subseteq q$. As $l(\mathbf{U}) \leq l(\mathbf{W})$ it follows $p[r + Kl(\mathbf{W})] \subseteq p[r + Kl(\mathbf{U})] \subseteq q$. By Lemma 7 and Lemma 9, we know $q[r + Kl(\mathbf{U})]$ extends to a tight multigeodesic. We may now apply Lemma 6 and Lemma 8 to find $q[(r + Kl(\mathbf{U})) + K] \subseteq q'$. It follows $p[r + Kl(\mathbf{W})] \subseteq q'$, completing the induction. \square

We now choose any tight multigeodesic z connecting a curve in $B(\alpha, r)$ to a curve in $B(\beta, r)$ and form a new multipath p by concatenating a tight multigeodesic connecting α to the end of z in $B(\alpha, r)$, the tight multigeodesic z , and a tight multigeodesic connecting the end of z in $B(\beta, r)$ to β . We take any word \mathbf{W} in \mathbf{S} and \mathbf{T} compatible with p and of maximal length among all such words. The number of \mathbf{S} operations appearing in the word \mathbf{W} is at most $4r$, for the multipath p has length at most $d(\alpha, \beta) + 4r$, no multipath connecting α and β has length strictly less than $d(\alpha, \beta)$, and length strictly decreases with each \mathbf{S} . Moreover, each \mathbf{S} operation will entail at most 4 tightening operations \mathbf{T} , for by Lemma 3 upon each \mathbf{S} we are only ever required to tighten at the ends of the shortcutting multipath (contributing at most two \mathbf{T} operations) and then at the two vertices either side of the shortcutting multipath (contributing at most two more \mathbf{T} operations). There are thus at most $4r \times 4$, or $16r$, \mathbf{T} operations in any such word \mathbf{W} . We conclude that the length of \mathbf{W} is at most $4r + 16r$, or $20r$.

Let us denote by q any multipath such that $p\mathbf{W}q$. As \mathbf{W} is of maximal length, q is both tight and 3-embedded. Lemma 10 tells us every curve belonging to a vertex of z and at distance strictly greater than $r + Kl(\mathbf{W})$, or Cr , from both α and β also belongs to q . The bound offered by Theorem 5 for the vertices of q thus applies to all the corresponding vertices of z . This concludes the proof of Lemma 4. \square

References

- [1] M. Bestvina, K. Fujiwara, *Bounded cohomology of subgroups of the mapping class groups* : Geom. Topol. **6** (2002) 69–89.
- [2] B. H. Bowditch, *Tight geodesics in the curve complex* : Invent. Math. **171** (2008) 281–300.
- [3] J. F. Brock, R. D. Canary, Y. N. Minsky, *The classification of Kleinian surface groups, II: The Ending Lamination Conjecture* : arXiv:math/0412006v1.
- [4] K. Fujiwara, *Subgroups generated by two pseudo-Anosov elements in a mapping class group. I. Uniform exponential growth* : in “Groups of Diffeomorphisms” ASPM **52**, Math. Soc. of Japan (2008) 283–296.

- [5] E. Ghys, P. de la Harpe, *Sur les groupes hyperboliques d'après Mikhael Gromov* : Birkhäuser (1990).
- [6] J. L. Harer, *The virtual cohomological dimension of the mapping class groups of an orientable surface* : Invent. Math. **84** (1986) 157-176.
- [7] W. J. Harvey, *Boundary structure of the modular group* : in "Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference" (ed. I. Kra, B. Maskit), Ann. of Math. Stud. No. 97, P. U. P. (1981) 245-251.
- [8] N. V. Ivanov, *Mapping class groups* : in "Handbook of geometric topology" (ed. R. Daverman and R. Sher), Elsevier (2001) 523-633.
- [9] T. Kobayashi, *Heights of simple loops and pseudo-Anosov homeomorphisms* : in "Braids (Santa Cruz, CA, 1986)" Contemp. Math. **78** (1988) 327-338.
- [10] H. A. Masur, Y. N. Minsky, *Geometry of the complex of curves I: Hyperbolicity* : Invent. Math. **138** (1999) 103-149.
- [11] H. A. Masur, Y. N. Minsky, *Geometry of the complex of curves II: Hierarchical structure* : Geom. Func. Anal. **10** No. 4 (2000) 902-974.
- [12] Y. N. Minsky, *The classification of Kleinian surface groups, I: Models and bounds* : arXiv:math/0302208v3.
- [13] K. J. Shackleton, *Tightness and computing distances in the curve complex* : preprint.
- [14] K. J. Shackleton, *Aspects of the curve complex and the mapping class group* : University of Southampton PhD thesis (2006).

UNIVERSITY OF TOKYO IPMU
<http://member.ipmu.jp/kenneth.shackleton/>
kenneth.shackleton@ipmu.jp
kjs2006@alumni.soton.ac.uk