

A computational acylindricity theorem for the mapping class group.

Kenneth J. Shackleton

kjs2006@alumni.soton.ac.uk

<http://www.is.titech.ac.jp/~kjshack5>

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.

KEYWORDS: surface; curve complex; mapping class group.

2000 MSC: 57M50 (primary); 30F60 (secondary).

§0. Introduction.

Let Σ be a compact, connected and orientable surface of genus $g(\Sigma)$ and $|\partial(\Sigma)|$ boundary components. In [Harv], Harvey associates to Σ a simplicial complex, $\mathcal{C} = \mathcal{C}(\Sigma)$, 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.

The curve complex is constructed as follows. We shall say that a simple loop on Σ is *trivial* only if it bounds a disc and *peripheral* only if it bounds an annulus containing one component of $\partial\Sigma$. Let $X = X(\Sigma)$ 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 can be realised disjointly.

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 dimension $\kappa = \kappa(\Sigma)$ of $\mathcal{C}(\Sigma)$ is equal to $3g(\Sigma) + |\partial\Sigma| - 4$. We see that Σ is in fact non-exceptional if and only if $\kappa(\Sigma)$ is positive.

When this is the case, the curve complex can be endowed with a 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, \mathcal{G} , of the curve complex with which \mathcal{C} is quasi-isometric and where distances between vertices are integers. We refer to this simplicial graph as the *curve graph*. All distances will be taken in this graph, whose own path-metric we denote by d . It is to be noted that the curve graph is nowhere locally finite and, when $\kappa(\Sigma) \geq 2$, between any two vertices of distance at least 2 there always 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 homotopy. This group has a natural cocompact action on the curve graph and this has been exploited by numerous authors in numerous ways; see [Hare] or [I] 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 [BeF] who, prompted by an argument

of Luo itself an adaptation of an argument of Kobayashi, 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 [Bo] 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 much stronger than WPD, where the number of such mapping classes is only assumed to be finite and where one end of the geodesic is assumed to be 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 action on a hyperbolic metric space necessarily contain free groups of rank 2 (and thus of arbitrarily high rank); see [F].

At some stage, both proofs from [BeF] and [Bo] 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. Our main result may be viewed as a computational alternative to the acylindricity theorem of Bowditch.

Theorem 1 *There is a computable function $F: \mathbb{N}^4 \rightarrow \mathbb{N}$ such that the following holds. Suppose Σ is non-exceptional. Let r be any non-negative integer. Then, for any two curves α and β with $d(\alpha, \beta) \geq 70r + 5$, 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 $F(\iota(\alpha, \beta), r, g(\Sigma), |\partial\Sigma|)$.*

We remark the lower bound $70r + 5$ can most likely be improved. Nevertheless, this is stronger than the WPD property found by Bestvina-Fujiwara, and logically independent of Bowditch’s acylindricity theorem. While the bounding function, F , 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 would be to find an argument that overcomes this trade-off, yielding both computability and uniformity simultaneously.

ACKNOWLEDGEMENTS: The author wishes to thank the Centre Interfacultaire Bernoulli, EPFL, for its warm and generous hospitality. He also wishes to thank Brian Bowditch for many interesting and helpful conversations.

§1. Background and definitions.

Before starting our proof of Theorem 1 in §2 we recall a few definitions, the most important of which is that of a tight multigeodesic after Masur-Minsky.

§1.1 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 to be defined additively.

§1.2 Paths and multipaths. A *path* in the curve graph is a sequence of vertices $(\gamma_0, \dots, \gamma_n)$ such that, for each index $i \in \{0, \dots, n-1\}$, the curves γ_i and γ_{i+1} are distinct and have zero intersection number, thus spanning an edge. A *geodesic* in the curve graph is then a path whose length is precisely the distance between its ends.

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 shall refer to each multicurve ν_i as a *vertex* of the multipath, following the language of Masur-Minsky. 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 $i \in \{0, \dots, n\}$.

A multipath (ν_0, \dots, ν_n) is said to be *k-embedded* only 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 $\gamma_j \in \nu_j$. Finally, a multipath (ν_0, \dots, ν_n) is said to be *k-quasigeodesic* only if $n \leq d(\nu_0, \nu_n) + k$ for each curve $\gamma_0 \in \nu_0$ and each curve $\gamma_n \in \nu_n$.

§1.3 Tight multigeodesics. The notion of a tight multigeodesic was introduced by Masur-Minsky [MaMi] to address the lack of local finiteness in the curve graph. Though there always exist infinitely many geodesics connecting any pair of vertices of distance at least 2, whenever $\kappa(\Sigma) \geq 3$, Corollary 6.4 of [MaMi] states that the number of tight multigeodesics connecting any given pair of vertices is always finite. A slightly more general definition was later offered by Bowditch in [Bo], where the finiteness result of Masur-Minsky is strengthened. In [S], the author offers computable bounds on the number of tight multigeodesics, taking either definition, connecting any pair of vertices.

In this paper, we shall 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 in general position 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 boundary of the resulting subsurface 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 . An example of a relative boundary is given below in Figure 1.




Figure 1: The multicurve ν_1 is the relative boundary of the curves ν_0 and ν_2 .

We are now ready to recall the relevant definition of a tight multigeodesic.

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

Let us now suppose that $(\nu_0, \nu_1, \nu_2, \nu_3)$ 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 [MaMi] 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 is a fundamental observation, a proof of which is implicit in the fourth paragraph of the proof of Lemma 4.5 from [MaMi].

As we shall be extensively referring to this tightening procedure let us summarise the above discussion with two statements, the second of which implies the first.

Lemma 3 (MaMi) *Let $(\nu_0, \nu_1, \nu_2, \nu_3)$ be a multigeodesic tight at ν_2 , and let ν'_1 denote the relative boundary $\partial(\nu_0, \nu_2)$ of ν_0 and ν_2 . Then, $(\nu_0, \nu'_1, \nu_2, \nu_3)$ is a tight multigeodesic.*

Lemma 4 (MaMi) *Let $(\nu_0, \nu_1, \nu_2, \nu_3)$ be a multigeodesic tight at ν_2 , and let ν'_0 denote any multicurve such that $(\nu'_0, \nu_1, \nu_2, \nu_3)$ is a multigeodesic. Let ν'_1 denote the relative boundary $\partial(\nu'_0, \nu_2)$ of ν'_0 and ν_2 . Then, $(\nu'_0, \nu'_1, \nu_2, \nu_3)$ is a tight multigeodesic.*

The tightening procedure is therefore 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 [MaMi] it is conceivable that tightening in a different order may produce distinct tight multigeodesics of the same ends, though this author is not aware of any such examples existing in the literature.

As we shall see in this paper, the tightening procedure can 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 only real catch is the multipath that results from tightening a 3-embedded multipath at a single vertex might not be 3-embedded. This, however, presents an obvious place 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 can remain intact if it happens to be geodesic over much of its length. Pertinent examples include multipaths obtained as the concatenation of a “short” multigeodesic, a “long” multigeodesic, and then another short multigeodesic. The full details of this procedure amount to the proof of Lemma 5 in §2.

§2. Proof of Theorem 1.

Let us introduce the notation $T(\alpha, \beta)$ for the set of all curves belonging to a tight multigeodesic 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, which may hold independent interest, 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. We momentarily postpone its proof.

Lemma 5 *There exists a computable function $F_1 : \mathbb{N}^4 \rightarrow \mathbb{N}$ such that the following holds. Let r be any non-negative integer and let α and β be any two curves. Then,*

$$|T(\alpha, \beta; r, 30r)| \leq F_1(\iota(\alpha, \beta), r, g(\Sigma), |\partial\Sigma|).$$

Given this, a proof of Theorem 1 is soon completed.

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 $T(\alpha, \beta; r, 30r)$. According to Lemma 5, the number of such translates of either δ or η is at most $F_1(\iota(\alpha, \beta), g(\Sigma), |\partial\Sigma|)$. Combining this fact with Lemma 7.4 from [Bo], 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 $\kappa(\Sigma)$, we complete a proof of Theorem 1. \diamond

To prove Lemma 5 we shall rely on the following theorem, a proof of which is implicit in the proof of Theorem 6 from [S]. The argument found therein is by contradiction, and constructs a path of length at most 2 between two curves of distance at least 3 and as such needs only the 3-embedded property of tight multigeodesics to apply.

Theorem 6 (S) *There exists a computable function $F_2 : \mathbb{N}^4 \rightarrow \mathbb{N}$ such that the following holds. Suppose Σ is non-exceptional, and let k denote any non-negative integer. Then, for any k -quasigeodesic, 3-embedded and tight multipath (ν_0, \dots, ν_n) we have, for each index j ,*

$$\iota(\nu_0, \nu_j) + \iota(\nu_j, \nu_n) \leq F_2(\iota(\nu_0, \nu_n), k, g(\Sigma), |\partial\Sigma|).$$

Proof: [of Lemma 5] For any multigeodesic z beginning in $B(\alpha, r)$ and ending in $B(\beta, r)$, suppose there exists a tight $2r$ -quasigeodesic multipath q connecting α to β , containing every curve from z at least distance $30r$ from both α and β , and which is 3-embedded. According to Theorem 6, each vertex of q has intersection number with α and intersection number with β uniformly bounded in terms of the quasigeodesic parameter of q , and hence in terms of r , and $\iota(\alpha, \beta)$, $g(\Sigma)$ and $|\partial\Sigma|$. Each vertex of z at least distance $30r$ from both α and β thus has similarly bounded intersection number with α and with β . This is enough to explicitly bound the cardinality of $T(\alpha, \beta; r, 30r)$.

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 relations between multipaths, denoted \rightarrow and \Rightarrow .

The relation \rightarrow will describe the shortening of a given multipath when it fails to be 3-embedded. The relation \Rightarrow will describe the tightening of a 3-embedded multipath at a single vertex. When the original multipath is geodesic over much of its length these operations are localised and much of the original multipath remains intact. In particular, this is the case when our multipath is the concatenation of a short multigeodesic, a long multigeodesic, and then another short multigeodesic each ending on single curves.

The interaction of \rightarrow and \Rightarrow need not be entirely well-behaved for the good work done by a sequence of \rightarrow might be undone by a single \Rightarrow . That is 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 in the direction of \rightarrow . In addition, this behaviour is prevented from seeping into the middle of our multipaths by Lemma 4. Thus our tightening procedure will stabilise, at which point we have found a tight and 3-embedded multipath connecting α to β and containing much of the original tight multigeodesic z , as required.

To be more precise, for two multipaths p and p' with common ends we shall write $p \rightarrow p'$ only if there exist consecutive vertices (μ_0, \dots, μ_i) of p , with $i \geq 2$, and consecutive vertices $(\omega_0, \dots, \omega_j)$ of p' , with $j \leq \min\{2, i - 1\}$, such that:

- $(\omega_0, \dots, \omega_j)$ is a tight multigeodesic;
- $p - \{\mu_0, \dots, \mu_i\} = p' - \{\omega_0, \dots, \omega_j\}$;
- $\omega_0 \subseteq \mu_0$, and
- $\omega_j \subseteq \mu_i$.

Note that the length of p is strictly greater than the length of p' , and that p' may fail to be tight at the two vertices of p' either side of the shortcut $(\omega_0, \dots, \omega_j)$ and its ends. An instance of the relation \rightarrow is depicted below in Figure 2.

For a 3-embedded multipath q and a multipath q' with the same ends but distinct from q , we shall write $q \Rightarrow q'$ only if q' results from tightening q at a single vertex. Notice even if q is 3-embedded and $q \Rightarrow q'$ we do not presume q' to be 3-embedded. If $q \Rightarrow q'$, then q and q' have the same length and do agree on all but one vertex.




Figure 2: The multipath p' is formed by shortcutting the multipath p .

Choose any tight multi-geodesic 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 multi-geodesic connecting α to the end of z in $B(\alpha, r)$, the tight multi-geodesic z , and a tight multi-geodesic connecting the end of z in $B(\beta, r)$ to β . Now take any sequence

$$p = p_0^0 \rightarrow p_1^0 \rightarrow \cdots \rightarrow p_{k_0}^0 \Rightarrow p_0^1 \rightarrow p_1^1 \rightarrow \cdots \rightarrow p_{k_t}^t \Rightarrow p_0^{t+1},$$

or $\cdots \rightarrow p_{k_t}^t$ if $p_{k_t}^t$ is already tight, that is of maximal length and such that $p_{k_i}^i$ is 3-embedded for each i . Note the number of \rightarrow relations appearing in the sequence is at most $2r$, for path length strictly decreases in the direction of \rightarrow . Moreover, the number of \Rightarrow in the sequence is at most $8r$ since, by Lemma 4, for each \rightarrow we are only ever required to tighten at the ends of the shortcutting multipath (contributing at most $4r$) and then at the two vertices either side of the shortcutting multipath (contributing at most $4r$). We conclude that the length of the sequence is at most $2r + 8r$, or $10r$.

Let us denote the final multipath in our sequence by q . By the maximality of our sequence, q is tight and 3-embedded. Since p can only fail to be 3-embedded inside the two balls of radius $r + 2$ about either end of z , so every curve belonging to z and at distance at least $30r$ from both α and β also belongs to q . With this, we conclude a proof of Lemma 5. \diamond

References.

- [BeF] M. Bestvina, K. Fujiwara, *Bounded cohomology of subgroups of the mapping class groups* : Geometry & Topology **6** (2002) 69 - 89.
- [Bo] B. H. Bowditch, *Tight geodesics in the curve complex* : to appear in Invent. Math.
- [F] K. Fujiwara, *Subgroups generated by two elements* : preprint (2007).
- [Hare] J. L. Harer, *The virtual cohomological dimension of the mapping class groups of an orientable surface* : Invent. Math. **84** (1986) 157-176.
- [Harv] 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), Annals of Mathematical Studies No. 97, Princeton University Press (1981) 245-251.
- [I] N. V. Ivanov, *Mapping class groups* : in “Handbook of geometric topology” (ed. R. Daverman and R. Sher), Elsevier (2001) 523-633.
- [MaMi] H. A. Masur, Y. N. Minsky, *Geometry of the complex of curves II: Hierarchical structure* : Geom. Funct. Anal. **10** (2000) 902-974.
- [S] K. J. Shackleton, *Tightness and computing distances in the curve complex* : preprint, Southampton (2004).

Kenneth J. Shackleton
Department of Mathematical and Computing Sciences
Tokyo Institute of Technology
2-12-1 O-okayama
Meguro-ku
Tokyo 152-8552
Japan
e-mail: shackleton.k.aa@m.titech.ac.jp
e-mail: kjs2006@alumni.soton.ac.uk
URL: <http://www.is.titech.ac.jp/~kjshack5>