

Tightness and computing distances in the curve complex

KENNETH J. SHACKLETON

We address the lack of local-finiteness in Harvey’s curve complex by computational means, computing bounds on certain intersection numbers among curves lying on a natural family of geodesics. We give a finite time algorithm for constructing all of these geodesics between any two curves, and treat stable lengths.

[57M50](#); [05C12](#)

1 Introduction

A borderless, connected and orientable finite type topological surface is *non-exceptional* if and only if it is not a sphere with at most four punctures or a torus with at most one puncture.

In (13) Harvey associates to such a surface Σ its curve complex, a simplicial complex whose 1-skeleton we call the *curve graph* and denote by $\mathcal{G}(\Sigma)$. By encoding some of the asymptotic geometry of Teichmüller’s metric, the curve graph plays a central role in the celebrated proof of Minsky and his collaborators of Thurston’s ending lamination conjecture (25; 5). A key step is a theorem of Masur-Minsky’s (23; 3) that every curve graph is hyperbolic in the sense of Gromov. However, the curve graph is nowhere locally finite and the theory of locally compact hyperbolic spaces does not apply.

To overcome this issue Masur-Minsky (24) identified the *tight geodesics*, a natural class of geodesic paths invariant under the action of the mapping class group, and proved any two vertices of the curve graph are connected by at least one and only finitely many tight geodesics. The arguments of Masur-Minsky (24) and of Bowditch (4) are not constructive as they involve taking limits of sequences, and the argument of Bowditch yields uniform but non-computed bounds on the cardinality of “slices”. We propose a constructive alternative offering computable but non-uniform bounds.

Theorem 1.1 *Let Σ be any non-exceptional surface. There exists a computable function $F : \mathbb{N} \rightarrow \mathbb{N}$ such that, for any tight geodesic $(\alpha_0, \dots, \alpha_n)$ in $\mathcal{G}(\Sigma)$, both $\iota(\alpha_0, \alpha_j)$ and $\iota(\alpha_j, \alpha_n)$ are at most $F(\iota(\alpha_0, \alpha_n))$ for each index j .*

From this we recover Masur-Minsky’s finiteness theorem.

Theorem 1.2 (Masur-Minsky) *Let Σ be any non-exceptional surface. Between any two vertices of $\mathcal{G}(\Sigma)$ there are only finitely many tight multigeodesics.*

Geodesic paths can be regarded as examples of multigeodesics, as indicated in Section 2.15. A “multigeodesic” is a sequence of sets of pairwise disjoint curves such that successively choosing a single curve arbitrarily from each always determines a geodesic path. See Section 2 for definitions.

In Section 3 we concentrate on proving the following key statement.

Proposition 1.3 *Let Σ be any non-exceptional surface. There exists a computable increasing function $F_1 : \mathbb{N} \rightarrow \mathbb{N}$ such that $k \leq F_1(k)$ for any non-negative integer $k \in \mathbb{N}$ and such that, for any multigeodesic $(\omega_0, \dots, \omega_n)$ in $\mathcal{G}(\Sigma)$ of length at least 2 and tight at ω_1 , we have $\iota(\omega_1, \omega_n) \leq F_1(\iota(\omega_0, \omega_n))$.*

An inductive argument given in Section 4 yields the following, and Theorem 1.1.

Corollary 1.4 *Let Σ be any non-exceptional surface. There exists a computable increasing function $F : \mathbb{N} \rightarrow \mathbb{N}$ such that for any tight multigeodesic $(\omega_0, \dots, \omega_n)$ in $\mathcal{G}(\Sigma)$, both $\iota(\omega_0, \omega_j)$ and $\iota(\omega_j, \omega_n)$ are at most $F(\iota(\omega_0, \omega_n))$ for each index j .*

The curves found on tight geodesics connecting two given curves belong to a computable set whose cardinality can be given an explicit upper bound. From this we can construct in explicitly bounded time all of the tight geodesics connecting any two curves.

Theorem 1.5 *There exists a finite time algorithm which takes as input a non-exceptional surface Σ and two curves α and β on Σ , and returns all the tight geodesics in $\mathcal{G}(\Sigma)$ between α and β .*

The two curves may be given as cyclic words in some finite generating set of the fundamental group of the surface. That said, the algorithm is admittedly impractical, with astronomical running time even for fairly small distances. Nevertheless, reading off the length of any geodesic path found this way does prove the following statement.

Corollary 1.6 *There exists a finite time algorithm which takes as input a non-exceptional surface Σ and two curves α and β on Σ , and returns $d(\alpha, \beta)$.*

Remark A version of Corollary 1.6 for closed surfaces of genus at least 2 is given in the unpublished thesis of Jason Leasure, see Corollary 3.2.6 of (19) where Leasure finds in a finite search space of curves a geodesic connecting two given curves. While Leasure treats closed surfaces it is likely, given the similarities that exist between our methods here, that his work can be adapted to treat the case of surfaces with finitely many punctures. Leasure’s work was brought to our attention by Richard P. Kent IV (17), by which time we had unwittingly submitted a manuscript for publication.

Remark In giving a proof to Proposition 1.3 it will transpire that assuming our multipaths are geodesic along their whole length is far stronger than we need. We do not exploit this here, instead exploring such matters in (28) to study the action of the mapping class group on the curve graph from an entirely computational perspective.

Acknowledgements

The author wishes to thank Brian Bowditch, Eric Katerman, Kasra Rafi and Makoto Sakuma, and to thank Max Planck Institut für Mathematik Bonn for its hospitality while writing up. The author also wishes to thank the referee for making many useful comments. The author was financially supported by an EPSRC PhD studentship and is grateful for this support.

2 Background

We supply all the definitions and background needed both to understand the statements of our results and their proofs.

2.1 Surfaces

A *surface* Σ is a borderless, connected and orientable topological 2-manifold with finitely generated fundamental group, of genus $g(\Sigma)$ with $n(\Sigma)$ punctures. Surfaces are uniquely determined up to homeomorphism by their genus and number of punctures (27; 26). We define $\xi(\Sigma) := \max\{3g(\Sigma) + n(\Sigma) - 3, 0\}$, a quantity commonly referred to as the *complexity* of Σ . The surface is *exceptional* if and only if $\xi(\Sigma) \leq 1$ and *non-exceptional* otherwise. We sometimes speak of a subset of a surface, such as a circle or an interval, as “being on” the surface.

2.2 Homotopy versus isotopy

Two subsets C_0 and C_1 of Σ are *freely homotopic* if there exists a continuous map $H : C_0 \times [0, 1] \rightarrow \Sigma$ such that $H(C_0 \times \{i\}) = C_i$ for $i \in \{0, 1\}$ and *freely isotopic* if in addition the homotopy H is an *isotopy*, so that $H|_{C_0 \times \{t\}} : C_0 \times \{t\} \rightarrow \Sigma$ is a homeomorphism onto its image for all $t \in [0, 1]$. Given a third subset B , the sets C_0 and C_1 are *freely homotopic* (resp. *freely isotopic*) *relative to B* if there exists a homotopy (resp. isotopy) H such that $H^{-1}(B) = (C_0 \cap B) \times [0, 1]$.

2.3 Loops and curves

A *loop* on a surface is a homeomorphic image of the circle $\mathbb{S}^1 = \{z \in \mathbb{C} : |z| = 1\}$. Thus by definition all loops are simple and unoriented.

A loop is said to be *trivial* if it is the image of the restriction to \mathbb{S}^1 of a continuous map from the closed disc $\mathbb{D} = \{z \in \mathbb{C} : |z| \leq 1\}$. A loop is said to be *peripheral* if it is the image of the restriction to \mathbb{S}^1 of a continuous map from the punctured disc $\mathbb{D} \setminus \{0\}$ whose image is not contained in any compact set. A loop is said to be *essential* if it is neither trivial nor peripheral.

On a surface two essential loops are freely homotopic if and only if they are freely isotopic (see (20), Proposition B.4.6 of (1) or Proposition 1.7 of (7)). We denote the free homotopy equivalence class of an essential loop c by $[c]$ and say that c *represents* the class $[c]$. Each such class is called a *curve*. We denote the set of all curves on Σ by $X(\Sigma)$. A surface Σ is non-exceptional if and only if $X(\Sigma)$ is not empty.

2.4 Intersection numbers

Associated to any pair of curves α and β is their *intersection number* $\iota(\alpha, \beta)$ defined equal to $\min\{|a \cap b| : a \in \alpha, b \in \beta\}$. Any curve can be represented by two disjoint loops and thus has zero intersection number with itself. A pair of curves of intersection number zero is said to be *disjoint* and is otherwise said to *intersect*.

We recall the following fact concerning intersection numbers.

Lemma 2.1 *Let Σ be any non-exceptional surface and let a and b be two essential loops on Σ such that $a \cap b$ is a finite set. Then $|a \cap b| = \iota([a], [b])$ if and only if there exists no closed disc $D \subset \Sigma$ such that ∂D is the union of a compact interval contained in a and a compact interval contained in b .*

A proof of Lemma 2.1 can be found in Section 1.2 of (7).

2.5 Transverse intersection

Given a pair of subsets a and b of a surface each homeomorphic to \mathbb{S}^1 or to the unit interval $[0, 1]$ such that $a \cap b$ is a finite set and given a common point $z \in a \cap b$, we say a and b *intersect transversely at z* if for every open set U containing z there exists an open regular neighbourhood V with $z \in V \subseteq U$ such that $a \cap V$ and $b \cap V$ are both homeomorphic to the real line and such that any path in V connecting any two distinct points of $(a \cap V) \setminus \{z\}$ separated by z along $a \cap V$ has non-empty intersection with b .

2.6 Multiloops and multicurves

A *multiloop* is a non-empty union of pairwise disjoint and pairwise non-homotopic essential loops. A *multicurve* is a non-empty set of pairwise distinct and pairwise disjoint curves, also viewed as the free homotopy class of a multiloop. We may view a loop as a multiloop with one component, and a curve as a multicurve with one element. We denote the free homotopy class of a multiloop ν by $[\nu]$. The intersection number of two multicurves or a curve and a multicurve is defined additively.

A *pants decomposition* is a multicurve that is maximal subject to inclusion. All pants decompositions of Σ have cardinality $\xi(\Sigma)$.

2.7 Hyperbolic geometry

The Klein-Poincaré uniformization theorem implies that a surface of negative Euler characteristic such as a non-exceptional surface admits a complete hyperbolic Riemannian metric of constant negative curvature -1 . Any curve (resp. multicurve) is uniquely represented by a simple closed geodesic (resp. union of pairwise disjoint simple closed geodesics) in such a metric. (See Lemma 2.3 of (6), Proposition B.4.7 of (1), Proposition 1.5 of (7) or Lemma 2.4.4 of (16) for instance.) Given two curves or multicurves μ and ω on Σ , if $u \in \mu$ and $v \in \omega$ denote their unique geodesic representatives in a hyperbolic metric then $\iota(\mu, \omega) = |u \cap v|$. (See Lemma 2.6 of (6) or Lemma 1.4 of (7) for instance.)

We recall the following fact concerning the lifts of two loops to the universal cover.

Lemma 2.2 *Let Σ be any non-exceptional surface and let a and b be two essential loops on Σ such that $|a \cap b| = \iota([a], [b])$. Then, for any universal covering map $\pi : \mathbb{H}^2 \rightarrow \Sigma$, the intersection of any component of the preimage $\pi^{-1}(a)$ and any component of the preimage $\pi^{-1}(b)$ is either empty or a single point.*

A proof of Lemma 2.2 can be found in Section 1.2 of (7).

2.8 Arcs

An *arc* on a surface is a homeomorphic image of the compact interval $[0, 1]$ of real numbers with the standard topology. All arcs are therefore simple and unoriented. On a surface two arcs with common endpoints are homotopic relative to their endpoints if and only if they are isotopic relative to their endpoints (9). An arc properly contained in a given arc, loop or multiloop ν is referred to as a *subarc* of ν .

2.9 Properly embedded arcs

For any subset $A \subseteq \Sigma$ we say that an arc h is *properly embedded in* (Σ, A) if $h \cap A = \partial h$ and there exists no closed disc whose boundary is the union of a compact subinterval of h and a compact interval contained in A .

More generally, for two subsets $A, B \subseteq \Sigma$ we say that an arc h is *properly embedded in* $(\Sigma, \{A, B\})$ if $\partial h \cap (A \cup B) = \partial h$, the intersection $h \cap A \cap B$ is empty, and there exists no closed disc whose boundary is the union of a compact subinterval of h and a compact interval entirely contained in at least one of A and B . An arc is properly embedded in (Σ, A) if and only if it is properly embedded in $(\Sigma, \{A, \emptyset\})$.

2.10 Parallel arcs

We say that two properly embedded arcs are *parallel in* (Σ, A) if they are homotopic relative to A in the standard sense. To be precise, two properly embedded arcs h_0 and h_1 are *parallel in* (Σ, A) if there exists a continuous map $H : h_0 \times [0, 1] \rightarrow \Sigma$ such that $H^{-1}(A) = (h_0 \cap A) \times [0, 1]$ and $H(h_0 \times \{i\}) = h_i$ for $i \in \{0, 1\}$.

More generally, two arcs h_0 and h_1 such that $\partial h_0 \subseteq A$ and $\partial h_1 \subseteq A$ are *parallel in* (Σ, A) if there exist two sets $\{h_0^1, \dots, h_0^n\}$ and $\{h_1^1, \dots, h_1^n\}$ of arcs properly embedded in (Σ, A) such that $h_0 = \bigcup_1^n h_0^i$ and $h_1 = \bigcup_1^n h_1^i$, such that both $h_0^i \cap h_0^{i+1}$ and $h_1^i \cap h_1^{i+1}$ are connected sets for each index $i \in \{1, \dots, n-1\}$, and such that h_0^i and h_1^i are parallel for each index $i \in \{1, \dots, n\}$.

If h_0 and h_1 are parallel in (Σ, A) and h_1 and h_2 are parallel in (Σ, A) then h_0 and h_2 are parallel in (Σ, A) .

2.11 Mapping class group

The *mapping class group* $\text{Map}(\Sigma)$ is defined as the group of all homeomorphisms of Σ modulo the normal subgroup of those homeomorphisms homotopic to the identity. The mapping class group acts naturally on $X(\Sigma)$ via the group action of homeomorphisms on the set of all loops. These actions have been studied by numerous authors in numerous ways; see (12; 2; 4; 28) and references contained therein.

A mapping class $\phi \in \text{Map}(\Sigma)$ is said to be *pseudo-Anosov* if no non-zero power of ϕ fixes a curve, equivalently if $\iota(\alpha, \phi^n \beta)$ grows exponentially in n for any two curves $\alpha, \beta \in X(\Sigma)$ (see (8)). A mapping class ϕ is said to be *partially pseudo-Anosov* if no non-zero power of ϕ fixes a pants decomposition, equivalently if $\iota(\mu, \phi^n \omega)$ grows exponentially in n for any two pants decompositions μ and ω .

2.12 Harvey's curve graph

In (13) Harvey associates to a surface Σ a flag simplicial complex of simplicial dimension $\xi(\Sigma) - 1$ called the curve complex. We study the 1-skeleton of this complex, denoted $\mathcal{G}(\Sigma)$ and referred to as the *curve graph*. This is defined independently by taking the vertex set to be $X(\Sigma)$ and declaring two vertices to be joined by an edge if and only if they are distinct and, as curves, they are disjoint.

If $\xi(\Sigma) \geq 2$ the curve graph is connected and can be endowed with the canonical path-metric d by realising $\mathcal{G}(\Sigma)$ as a length space with each edge of length 1. The distance between two curves is the length of any shortest path between them. Since the action of the mapping class group preserves intersection number, the action is by isometries. With only a few exceptions, the isometry group of the curve graph is in fact isomorphic to the mapping class group (15; 18; 22). For any two curves $\alpha, \beta \in X(\Sigma)$ and $\phi \in \text{Map}(\Sigma)$ a pseudo-Anosov mapping class, the distance $d(\alpha, \phi^n \beta)$ grows linearly in n (24).

The curve graphs all have infinite diameter (23; 14) but it is a simple matter to effectively characterise the first three distances: Two curves are at distance 0 if and only if they are equal; are at distance 1 if and only if they are distinct and disjoint, and are at distance 2 if and only if they have positive intersection number and there exists a curve disjoint from both. Two curves α and β are at distance at least 3 if and only if every curve has positive intersection number with at least one of α and β , whereupon $\{\alpha, \beta\}$ may be said to *fill* the surface. Given two loops a and b representing a pair of curves $a \in \alpha$ and $b \in \beta$ that fills the surface, the complement of the union $a \cup b$ is the disjoint union of finitely many open discs and once-punctured open discs.

The distance between a pair of curves and their intersection number are related in the following well-known lemma, recalled from Lemma 2.1 in (14) with the upper bound cast as an integer.

Lemma 2.3 *If $\xi(\Sigma) \geq 2$, we have $d(\alpha, \beta) \leq \lfloor 2\log_2(\iota(\alpha, \beta)) \rfloor + 2$ for all $\alpha, \beta \in X(\Sigma)$.*

Upper bounds on distance in terms of intersection number have been circulating for some time though the earliest in print appears to be Lemma 2.1 of (23) or Lemma 2.1 of (14). Their proofs typically use an argument given by Lickorish (21). In general one can do no better than logarithmic given the existence pseudo-Anosov mapping classes.

We say that two multicurves μ and ω are of *distance* n , and write $d(\mu, \omega) = n$, if $d(\gamma, \delta) = n$ for each $\gamma \in \mu$ and $\delta \in \omega$. For two multicurves μ and ω of distance $d(\mu, \omega) \geq 1$, if u and v are their respective geodesic representatives in a hyperbolic metric on the surface then $\iota(\mu, \omega) = |u \cap v| = |u \pitchfork v|$.

2.13 Intervals

A *decomposition* of a compact interval J into compact intervals is a set \mathcal{J} of compact intervals such that $J = \bigcup_{\mathcal{J}} I$ and such that, for all I_1 and I_2 in \mathcal{J} , the intersection $I_1 \cap I_2$ is infinite if and only if $I_1 = I_2$.

An *interval of integers* is a non-empty subset of \mathbb{Z} of the form $\{n \in \mathbb{Z} : i \leq n \leq j\}$ for some $i, j \in \mathbb{Z} \sqcup \{\pm\infty\}$ with $i \leq j$.

2.14 Paths and multipaths

We regard a *path* in the curve graph as a sequences of curves $(\gamma_i)_i$ indexed by an interval of integers such that consecutive curves are distinct and disjoint, equivalently they span an edge in the curve graph. A *multipath* is a sequence of multicurves $(\omega_i)_i$ indexed by an interval of integers such that, for each $\gamma_i \in \omega_i$ and each i , the sequence $(\gamma_i)_i$ is a path. If $(\gamma_i)_i$ is a path then $(\{\gamma_i\})_i$ is a multipath.

If a given path or multipath has a finite minimal or maximal index we refer to such an index and the corresponding curve or multicurve as *terminal*. We refer to all the other indices and their corresponding curves or multicurves as *non-terminal*.

2.15 Geodesics and multigeodesics

A *geodesic* is a path $(\gamma_i)_i$ such that $d(\gamma_i, \gamma_j) = |i - j|$ for any two finite indices i and j . A multipath $(\omega_i)_i$ is a *multigeodesic* if for each $\gamma_i \in \omega_i$ and each i the sequence $(\gamma_i)_i$ is a geodesic path. If $(\gamma_i)_i$ is a geodesic then $(\{\gamma_i\})_i$ is a multigeodesic.

2.16 Relative boundary

To any pair of multicurves ω_0 and ω_2 at distance 2, Masur-Minsky (24) associate a multicurve called its *relative boundary* $\partial(\omega_0, \omega_2)$ as follows. First, choose a representative multiloop v_0 of ω_0 and a representative multiloop v_2 of ω_2 such that $|v_0 \cap v_2| = \iota(\omega_0, \omega_2)$. Take an open set U that admits a neighbourhood deformation retraction onto the union $v_0 \cup v_2$ and whose boundary is a disjoint union of finitely many loops. Attach to U all the closed disc and once-punctured closed disc components of $\Sigma \setminus U$ and denote the resulting subsurface of Σ by U^* . The free homotopy class of any multiloop maximal subject to inclusion and contained in ∂U^* is a non-empty and well-defined multicurve associated to the pair (ω_0, ω_2) , its relative boundary.

2.17 Tight geodesics

Definition 2.4 (Masur-Minsky) A multigeodesic $(\omega_i)_i$ is *fully tight at* ω_j for some non-terminal index j if $\omega_j = \partial(\omega_{j-1}, \omega_{j+1})$. We say $(\omega_i)_i$ is *fully tight* if it is fully tight at each non-terminal ω_j .¹

The existence of a fully tight multigeodesic between any two curves is established in the proof of Lemma 4.5 from (24) as follows: Let us suppose $(\omega_0, \omega_1, \omega_2, \omega_3)$ is a multigeodesic fully tight at ω_2 . We replace ω_1 with the relative boundary $\partial(\omega_0, \omega_2)$ of ω_0 and ω_2 . It can be shown that this does not affect the tightness at ω_2 , that is $(\partial(\omega_0, \omega_2), \omega_2, \omega_3)$ is fully tight. This tightening operation is therefore stable and can be applied to the vertices of a geodesic path in any order to produce a fully tight multigeodesic with the same endpoints.

Between a pair of curves of distance 2 there is only one fully tight multigeodesic. In general it is conceivable that tightening a given geodesic or multigeodesic in a different order may produce a different multipath.

An important observation concerning the relative boundary is the following.

¹The term ‘‘tight multigeodesic’’ is used in (24), ‘‘fully tight’’ is the author’s own term.

Lemma 2.5 Suppose ω_0 and ω_2 are two multicurves at distance 2. If $\delta \in X(\Sigma)$ is any curve such that $\iota(\delta, \partial(\omega_0, \omega_2)) > 0$, then $\iota(\delta, \omega_0) > 0$ or $\iota(\delta, \omega_2) > 0$.

This brings us to a slight generalisation of Definition 2.4.

Definition 2.6 (Bowditch) A multigeodesic $(\omega_i)_i$ is *tight at* ω_j for some non-terminal index j if for any curve $\delta \in X(\Sigma)$ whenever $\iota(\delta, \omega_j) > 0$ we have $\iota(\delta, \omega_{j-1}) > 0$ or $\iota(\delta, \omega_{j+1}) > 0$. We say $(\omega_i)_i$ is *tight* if it is tight at each non-terminal ω_j .

We note that any fully tight multigeodesic is tight, and that if $(\omega_0, \omega_1, \omega_2)$ is a tight multigeodesic then $\omega_1 \subseteq \partial(\omega_0, \omega_2)$.

Definition 2.7 (Bowditch) A geodesic $(\alpha_i)_i$ is *tight* if there exists a tight multigeodesic $(\omega_i)_i$ such that $\alpha_i \in \omega_i$ for each index i .

We note that in the curve graph of the five-times punctured sphere and the twice-punctured torus every geodesic path is tight for the topology of the two surfaces leaves no alternative.

We rely on Definition 2.6 only in the proof of Lemma 3.2.

2.18 Stable lengths

Given a metric space (X, d) and an isometry $\phi : X \rightarrow X$, the *stable length* $\|\phi\|$ of ϕ is defined equal to $\lim_{n \rightarrow \infty} d(x, \phi^n x)/n$, for any $x \in X$. It can be verified that $\|\phi\|$ is always finite and, by twice applying the metric triangle inequality, that $\|\phi\|$ does not depend on the choice of x .

3 Proof of Proposition 1.3

3.1 Curve surgery

We first take a pair of intersecting curves and return a new curve with useful properties.

Lemma 3.1 Let γ and β be any two curves on Σ , and let $c \in \gamma, b \in \beta$ be loops such that $|c \cap b| = \iota(\gamma, \beta)$. Let J be any compact non-empty subinterval of b . Suppose there exists a subarc h of c such that $|h \cap J| = 3$. Then, there exists a curve $\delta \in X(\Sigma)$ represented by a loop in every open subset of Σ containing $h \cup J$ such that $\iota(\delta, \gamma) > 0$ and such that $\iota(\delta, \beta) > 0$.

Proof We construct a parameterised loop $p : [0, 1] \rightarrow \Sigma$ as follows. Let x_1, x_2 and x_3 denote the three points of $h \cap J$ where x_2 separates x_1 and x_3 along h , and let h_1 denote the subarc of h connecting x_1 and x_2 and let h_2 denote the subarc of h connecting x_2 and x_3 . The construction of p falls into three types (see the left half of Figure 1).

- (Type 1) If there exists $i \in \{1, 2\}$ such that h_i is incident to J from both sides of b we form p by traversing h_i and then the subinterval of J spanned by ∂h_i .
- (Type 2) Otherwise, if x_2 does not separate x_1 and x_3 along J we form p by traversing h from x_1 and then the subinterval of J spanned by x_3 and x_1 .
- (Type 3) If instead x_2 separates x_1 and x_3 along J we form p by traversing h_1 from x_1 to x_2 , J from x_2 to x_3 , h_2 from x_3 to x_2 , and then J from x_2 to x_1 .

We extend p to a periodic map $p : \mathbb{R} \rightarrow \Sigma$ such that $p(t + 1) = p(t)$ for all $t \in \mathbb{R}$. Let $\pi : \mathbb{H}^2 \rightarrow \Sigma$ denote a universal covering map and \tilde{p} any π -lift of p . We denote $\tilde{p}(i)$ by z_i , the π -lift of b containing z_i by \tilde{b}_i , and the π -lift of c containing z_i by \tilde{c}_i for each integer $i \in \mathbb{Z}$. The union of all \tilde{b}_i will typically be a proper subset of the preimage $\pi^{-1}(b)$, and similarly the union of all \tilde{c}_i will typically be a proper subset of $\pi^{-1}(c)$.

We note $(\tilde{b}_i)_i$ and $(\tilde{c}_i)_i$ are both sequences of uniform simple quasi-geodesics, and that \tilde{b}_{i-1} and \tilde{b}_{i+1} are separated in \mathbb{H}^2 by \tilde{b}_i and that \tilde{c}_{i-1} and \tilde{c}_{i+1} are separated in \mathbb{H}^2 by \tilde{c}_i for every $i \in \mathbb{Z}$ regardless of the type defining p . Since $|c \cap b| = \iota(\gamma, \beta)$, according to Lemma 2.2 we know \tilde{b}_i intersects \tilde{c}_j in only at most one point for each $i, j \in \mathbb{Z}$. In particular we note that $z_0 \neq z_2$, so $p(\mathbb{R})$ is not homotopically trivial, and that $\tilde{p}(\mathbb{R})$ intersects every \tilde{b}_i and every \tilde{c}_i in a non-empty connected set. (See the right half of Figure 1.)

Regardless of the type defining p there exists a positive real number $\epsilon_0 > 0$ such that, for any positive real number ϵ with $0 < \epsilon < \epsilon_0$, at least one component \tilde{e} of the boundary of the open ϵ -neighbourhood of the image $\tilde{p}(\mathbb{R})$ in \mathbb{H}^2 projects to a loop $\pi(\tilde{e})$ under π . We denote this loop by e . Since \tilde{e} is homotopic to $\tilde{p}(\mathbb{R})$ we note e is a not a trivial loop. Moreover, \tilde{e} intersects each \tilde{b}_i and each \tilde{c}_j in a single point. It follows e is also not peripheral and must be essential, defining a curve δ whose intersection number with γ and with β is positive. \square

3.2 Motivating example

We motivate our arguments by offering a proof of Proposition 1.3 for multigeodesic paths of length 3. In so doing, we demonstrate how Lemma 3.1 will be applied.

Let $\pi = (\omega_0, \omega_1, \omega_2, \omega_3)$ be a multigeodesic tight at ω_1 . There exist representative multiloops v_i of ω_i for $i \in \{0, 1, 2, 3\}$ such that $v_i \cap v_j \cap v_k = \emptyset$ for integers $0 \leq i < j < k \leq 3$ and such that $\iota(\omega_i, \omega_j) = |v_i \cap v_j|$ for integers $0 \leq i < j \leq 3$. In particular and by assumption, $v_2 \cap v_3 = \emptyset$.

Suppose for contradiction that $\iota(\omega_1, \omega_3) \geq 2\kappa(\Sigma)\iota(\omega_0, \omega_3) + 1$. The pigeonhole principle tells us that there exists a non-empty compact interval $J \subset v_3 \setminus v_0$ such that $|v_1 \cap J| \geq 2\kappa(\Sigma) + 1$ and that there exists a loop component c_1 of v_1 such that $|c_1 \cap J| \geq 3$. Let h denote any subarc of c_1 such that $|h \cap J| = 3$.

Applying Lemma 3.1 to h along J we know there exists a curve δ represented by a loop in every open set containing $h \cup J$ such that $\iota(\delta, \omega_1) > 0$. Since π is tight at ω_1 the curve δ must also have positive intersection number with at least one of ω_0 and ω_2 . Now both h and J are disjoint from v_0 and so $\iota(\delta, \omega_0) = 0$. Thus, $\iota(\delta, \omega_2) > 0$ is the only possibility and it follows that $(h \cup J) \cap v_2 \neq \emptyset$. However, h is a subarc of v_1 and is therefore disjoint from v_2 . We conclude that $J \cap v_2 \neq \emptyset$ and as such $v_2 \cap v_3 \neq \emptyset$. This is absurd, and we complete the contradiction.

We have just proven that $\iota(\omega_1, \omega_3) \leq 2\kappa(\Sigma)\iota(\omega_0, \omega_3)$. QED

3.3 Realising curves

From now on $(\omega_0, \dots, \omega_n)$ is a multigeodesic tight at ω_1 and the multiloops v_i representing ω_i , for $i \in \{0, \dots, n\}$, together satisfy the following conditions:

- (1) the components of v_i are pairwise disjoint for each i ;
- (2) v_i and v_j have only transverse intersection for $i \leq j - 2$;
- (3) $v_i \cap v_{i+1} = \emptyset$ for each $i \leq n - 1$;
- (4) $|v_i \cap v_n| = \iota(\omega_i, \omega_n)$ for each $i \leq n - 2$, and
- (5) $v_i \cap v_j \cap v_n = \emptyset$ for $i + 1 < j < n - 1$.

We could for instance endow Σ with a fixed complete hyperbolic metric and represent each ω_i uniquely by a union of geodesics. This will simultaneously minimise all the pairwise intersection numbers but in particular ensure conditions (1)–(4) are satisfied. Perturbing these geodesics if need be, in addition we have (5). We draw the reader's attention to the distinction between the multiloop v_i and its multicurve $\omega_i = [v_i]$.

In what follows we denote by J an arbitrary non-empty compact subinterval of $v_n \setminus v_0$ and by c_n the component of v_n containing J .

3.4 Small intersection numbers

Lemma 3.2 *Suppose $|c_1 \cap J| \geq 3$ for some component c_1 of v_1 . Then $v_2 \cap J$ is non-empty.*

Proof Let h denote any subarc of c_1 such that $|h \cap J| = 3$. According to Lemma 3.1, there exists a curve δ represented by a loop in any open set containing $h \cup J$ and of positive intersection number with ω_1 . The tightness of the multigeodesic at ω_1 implies δ has positive intersection number with at least one of ω_0 and ω_2 . However v_0 is disjoint from v_1 and J , and v_2 is disjoint from v_1 . It follows $v_2 \cap J$ must be non-empty. \square

Corollary 3.3 *Suppose $|v_1 \cap J| \geq 2\xi(\Sigma) + 1$. Then $v_2 \cap J$ is non-empty.*

Proof There exists a component c_1 of v_1 such that $|c_1 \cap J| \geq 3$. The claim now follows from Lemma 3.2. \square

3.5 Nested intersection

Lemma 3.4 *For $2 \leq m < n$ suppose there exists a continuous map $H : [0, 1] \times [0, 1] \rightarrow \Sigma$ such that:*

- (1) $v_0 \cap \text{im}(H) = \emptyset$;
- (2) $H(\{\frac{1}{2}\} \times [0, 1]) \subseteq v_m$, and
- (3) $H([0, 1] \times \{0, \frac{1}{2}, 1\}) \subseteq J$.

Let $J_0 \subseteq J$ be any compact subinterval of J containing the image $H(\{\frac{1}{2}\} \times \{0, \frac{1}{2}, 1\})$. Then, the set $v_{m+1} \cap J_0$ is non-empty.

Proof We note that the existence of such a subinterval J_0 is assured by condition (3).

Appealing to conditions (2) and (3) there exists a compact connected 1-submanifold h of $H(\{\frac{1}{2}\} \times [0, 1])$ such that $|h \cap J_0| = 3$. Let c_m denote the component of v_m containing h and whose existence is also assured by condition (2). We note that h is not a loop because otherwise it represents a curve simultaneously of distance at least 2 from every component of ω_0 and, by condition (1), disjoint from v_0 and this would be absurd. Thus, h is a proper subarc of c_m . (See Figure 2.)

We appeal to Lemma 3.1, with $\gamma = [c_m], \beta = [c_n], c = c_m, b = c_n$ and to h along J_0 , to find a curve δ represented by a loop in any open set containing $h \cup J_0$. Since $v_0 \cap h = \emptyset, v_0 \cap J_0 = \emptyset$ and the topology of a surface is regular there exists an open set U containing $h \cup J_0$ such that $U \cap v_0 = \emptyset$. For any loop $e_m \in \delta$ representing δ such that $e_m \subset U$ it follows $e_m \cap v_0 = \emptyset$. We conclude that $\iota(\delta, \omega_0) = 0$.

We now note that $c_m \cap v_{m+1} = \emptyset$ because $c_m \subseteq v_m$ and $v_m \cap v_{m+1} = \emptyset$. If in addition $v_{m+1} \cap J_0 = \emptyset$ then, after replacing both U and e_m if need be, e_m is disjoint from v_{m+1} . It follows that $\iota(\delta, \omega_{m+1}) = 0$. Moreover, since $\iota(\delta, \omega_0)$ is zero and $m \geq 2$, we note that δ is distinct from each component of ω_{m+1} . We therefore have a multipath $(\omega_0, \{\delta\}, \omega_{m+1})$ of length 2 and so $d(\omega_0, \omega_{m+1}) \leq 2$. This is absurd, since $d(\omega_0, \omega_{m+1}) = m+1 \geq 2+1 = 3 > 2$. We deduce that $v_{m+1} \cap J_0$ must be non-empty, and with this we complete a proof of Lemma 3.4. \square

3.6 Parallel arcs

Lemma 3.5 *Let s be a non-negative integer. There exists a computable constant $K_0(s)$ such that, for any compact interval I on Σ , in any set of at least $K_0(s)$ pairwise homotopically disjoint arcs properly embedded in (Σ, I) there are at least s arcs parallel in (Σ, I) .*

Proof A set of pairwise homotopically disjoint pairwise non-parallel arcs properly embedded in (Σ, I) has cardinality at most $6g(\Sigma) + 2n(\Sigma) - 3$. By the pigeonhole principle, we can define $K_0(s) := \max\{(6g(\Sigma) + 2n(\Sigma) - 3)(s - 1) + 1, 0\}$. \square

Corollary 3.6 *Let s be a non-negative integer. There exists a computable constant $K(s)$ such that in any set \mathcal{A} of at least $K(s)$ pairwise homotopically disjoint arcs each intersecting a compact interval I on Σ only three times and whose boundaries are contained in I there are at least s arcs parallel in (Σ, I) .*

Proof A fairly crude constant $K(s)$ can be obtained as follows. Suppose \mathcal{A} contains at least $K_0(K_0(s))$ such arcs. Each of these arcs can be expressed as the union of two arcs properly embedded in (Σ, I) . By Lemma 3.5, there exist $K_0(s)$ arcs in \mathcal{A} containing a properly embedded subarc parallel to a common arc. Among these arcs, Lemma 3.5 implies there exist s parallel arcs. We thus set $K(s) = K_0(K_0(s))$. \square

3.7 Large intersection

Lemma 3.7 *There exists a computable and monotonic increasing exponential function $G : \mathbb{N} \rightarrow \mathbb{N}$ such that the following holds. For m a non-negative integer strictly less than n , if $|v_1 \cap J| \geq G(m)$ then there exists a sequence of nested subintervals $J_1 \supset J_2 \supset \dots \supset J_m$ of J such that:*

- (1) $\partial J_i \subset v_i$ for each $i \in \{1, \dots, m\}$, and
- (2) $v_{m+1} \cap J_m \neq \emptyset$.

Proof We define $G(0) := 0$ and note by Corollary 3.3 that $G(1) := 2\xi(\Sigma) + 1$ satisfies our assertion for $m = 1$ when J_1 is defined to be the intersection of all compact subintervals of J containing a common set comprising $G(1)$ pairwise distinct points of $v_1 \cap J$. We regard $m = 1$ as the base case for an induction on m and henceforth $m \geq 2$.

Suppose inductively we have found $G(m)$. If $|v_1 \cap J| \geq 3K(3)G(m)$ we can decompose J into the union of at least $3K(3)$ compact subintervals J_1^i , for $i \in \{1, \dots, 3K(3)\}$, each intersecting v_1 at least $G(m)$ times and such that $\partial J_1^i \subset v_1$ for each index i and such that $J_1^i \cap J_1^j$ is finite for indices $i < j$. The inductive hypothesis implies the existence of a nested sequence of compact intervals $J_1^i \supset J_2^i \supset \dots \supset J_m^i$ such that $\partial J_j^i \subset v_j$ and $v_{m+1} \cap J_m^i \neq \emptyset$ for each index i and j .

Let \mathcal{A} denote the set of all subarcs h of v_1 such that $|h \cap J| = 3$, $\partial h \subset J$ and $\partial h \cap \partial J_1^i$ is non-empty for some index i . The set \mathcal{A} has cardinality at least $K(3)$. We apply Corollary 3.6 to this set \mathcal{A} and (Σ, J) to find a family of three or more pairwise parallel arcs belonging to \mathcal{A} .

Since $v_0 \cap J = \emptyset$, among all of these parallel arcs there exist two that can be related by a homotopy $H : [0, 1] \times [0, 1] \rightarrow \Sigma$ such that $v_0 \cap \text{im}(H) = \emptyset$ and such that $H([0, 1] \times \{0\}) = J_1^i$ for some index i and $H([0, 1] \times \{0, \frac{1}{2}, 1\}) \subseteq J$. Since $v_j \cap v_{j+1} = \emptyset$ for each $j \in \{1, \dots, m\}$ and $J_1^i \supset J_2^i \supset \dots \supset J_m^i$ we may choose H so that in addition

$$I_j := H\left(\left\{\frac{1}{2} \pm \frac{m+1-j}{2m}\right\} \times [0, 1]\right) \subseteq v_j$$

for each $j \in \{1, \dots, m+1\}$ and where the images I_j are pairwise disjoint. In other words, the homotopy H can be one ‘‘swept out’’ by $J_1^i \supset \dots \supset J_m^i$ so to speak.

Let \widehat{J}_j denote the intersection of all compact subintervals of J containing $I_j \cap J$ for each $j \in \{1, \dots, m+1\}$. We note that $\widehat{J}_1 \supset \widehat{J}_2 \supset \dots \supset \widehat{J}_{m+1}$ and that $\partial \widehat{J}_j \subset v_j$ for each

such j . We also note that $H(\{\frac{1}{2}\} \times \{0, \frac{1}{2}, 1\}) \subset \widehat{J}_1$, because $I_{m+1} = H(\{\frac{1}{2}\} \times [0, 1])$ and $H(\{\frac{1}{2}\} \times \{0, \frac{1}{2}, 1\}) \subset I_{m+1} \cap J \subset \widehat{J}_{m+1} \subseteq \widehat{J}_1$.

We can now apply Lemma 3.4 along $\widehat{J}_1 \subseteq J$ to deduce $v_{m+2} \cap \widehat{J}_{m+1} \neq \emptyset$. We therefore define $G(m+1) = 3K(3)G(m)$ and then take $\widehat{J}_1 \supset \widehat{J}_2 \supset \cdots \supset \widehat{J}_{m+1}$ as the promised nested sequence of intervals, completing the induction. \square

Remark We may take $G(0) = 0$ and $G(n) = (2\xi(\Sigma) + 1)(3K(3))^{n-1}$ for all $n \in \mathbb{N}_+$.

3.8 Intersection bounds

A proof of Proposition 1.3 can be completed as follows. We define a function $F_1 : \mathbb{N} \rightarrow \mathbb{N}$ by

$$F_1(n) := nG(\lfloor 2\log_2 n \rfloor)$$

for all $n \in \mathbb{N}_+$ and $F_1(0) := 0$. We note that $n \leq F_1(n)$ for all non-negative integers $n \in \mathbb{N}$. We also note that $\iota(\omega_0, \omega_n)$ is non-zero as $d(\omega_0, \omega_n) \geq 2$ by assumption.

Suppose for contradiction $\iota(\omega_1, \omega_n) > F_1(\iota(\omega_0, \omega_n))$, noting n is then at least 3. The intersection number $\iota(\omega_1, \omega_n)$ can be expressed as the sum of $|v_1 \cap J_i|$ over each component J_i of $v_n \setminus v_0$, of which there are exactly $\iota(\omega_0, \omega_n)$, and so there exists a compact subinterval J of $v_n \setminus v_0$ such that

$$\iota(\omega_0, \omega_n)|v_1 \cap J| \geq \iota(\omega_1, \omega_n).$$

Combining the two inequalities, we find

$$\iota(\omega_0, \omega_n)|v_1 \cap J| > F_1(\iota(\omega_0, \omega_n)).$$

Substituting in the defining expression for F_1 yields

$$\iota(\omega_0, \omega_n)|v_1 \cap J| > \iota(\omega_0, \omega_n)G(\lfloor 2\log_2(\iota(\omega_0, \omega_n)) \rfloor)$$

and after dividing both sides by the non-zero integer $\iota(\omega_0, \omega_n)$ we have

$$|v_1 \cap J| > G(\lfloor 2\log_2(\iota(\omega_0, \omega_n)) \rfloor).$$

We recall $G(n)$ is monotonic increasing in n , and appealing to Lemma 2.3 we find

$$|v_1 \cap J| > G(n-2).$$

However, condition (2) of Lemma 3.7 implies $v_{n-1} \cap v_n$ is not empty and this is absurd. With this we complete the contradiction, and a proof of Proposition 1.3. QED

4 Proof of Corollary 1.4

Given a function $\Theta : \mathbb{N} \rightarrow \mathbb{N}$ we define a sequence of functions $(\Theta^m)_{m \in \mathbb{N}}$ inductively by first defining Θ^0 to be the identity on \mathbb{N} and then defining the iterate $\Theta^{m+1} := \Theta \circ \Theta^m$ for each non-negative integer $m \in \mathbb{N}$. We use the notation $\Theta^m(n)$ for the natural number obtained by evaluating the function Θ^m at n for each $m, n \in \mathbb{N}$. In general this will be distinct from the value of $\Theta(n)$ raised to the power of m , denoted $\Theta(n)^m$.

For each index $m \in \{1, \dots, n-2\}$, by Proposition 1.3 we have

$$\iota(\omega_m, \omega_n) \leq F_1(\iota(\omega_{m-1}, \omega_n)).$$

As $k \leq F_1(k)$ for all $k \in \mathbb{N}$, we have

$$F_1(\iota(\omega_{m-1}, \omega_n)) \leq F_1(F_1(\iota(\omega_{m-2}, \omega_n))) = F_1^2(\iota(\omega_{m-2}, \omega_n)) \leq \dots \leq F_1^m(\iota(\omega_0, \omega_n)).$$

If we define $F(n) = F_1^{\lfloor 2 \log_2(n) \rfloor}(n)$ for all $n \in \mathbb{N}$, then by Lemma 2.3 we have

$$\iota(\omega_m, \omega_n) \leq F(\iota(\omega_0, \omega_n))$$

for all indices m . By reversing the multi-geodesic and repeating this argument we find

$$\iota(\omega_0, \omega_m) \leq F(\iota(\omega_0, \omega_n))$$

for all indices m . This completes the proof of Corollary 1.4. QED

5 Proof of Theorem 1.2

Lemma 5.1 *Suppose Σ is a non-exceptional surface. Let $\alpha, \beta \in X(\Sigma)$ be two curves such that $d(\alpha, \beta) \geq 3$ and let K be any non-negative integer. Then, the set $\{\gamma \in X(\Sigma) : \iota(\alpha, \gamma) + \iota(\gamma, \beta) \leq K\}$ is computable and has finite cardinality uniformly and explicitly bounded above in terms of $\iota(\alpha, \beta)$, $n(\Sigma)$ and K .*

Proof Let $a \in \alpha$ and $b \in \beta$ be representative loops such that $|a \cap b| = \iota(\alpha, \beta)$. Since $d(\alpha, \beta) \geq 3$, the two curves together fill the surface. The complement of the union $a \cup b$ is therefore the disjoint union of finitely many open discs and once-punctured open discs. For any curve $\gamma \in X(\Sigma)$ we can find a loop $c \in \gamma$ such that $|c \cap a| = \iota(\gamma, \alpha)$, $|c \cap b| = \iota(\gamma, \beta)$ and $a \cap b \cap c = \emptyset$. We can thus express c as the union of at most $\iota(\alpha, \gamma) + \iota(\gamma, \beta)$ intervals each properly embedded in $(\Sigma, \{a, b\})$.

Let Λ denote the set of all loops c equal to the union of at most K compact intervals each properly embedded in $(\Sigma, \{a, b\})$. We say that two such loops c_0 and c_1 are *equivalent*

if there exists a continuous map $H : c_0 \times [0, 1] \rightarrow \Sigma$ such that $H(c_0 \times \{i\}) = c_i$ for $i \in \{0, 1\}$ and such that if $I \subset c_0$ is a compact interval properly embedded in $(\Sigma, \{a, b\})$ then $H(I \times \{t\})$ is similarly a properly embedded interval for all $t \in [0, 1]$. This defines an equivalence relation on the set Λ , and the set of equivalence classes is not only finite but also has cardinality that can be uniformly and explicitly bounded in terms of K and $n(\Sigma)$. Each equivalence class determines a curve and the number of curves that can be represented this way while using at most K properly embedded intervals is similarly bounded. Because every curve $\gamma \in X(\Sigma)$ such that $\iota(\alpha, \gamma) + \iota(\gamma, \beta) \leq K$ can be represented in this manner, we already have an upper bound on the number of such curves. \square

Consider any two curves $\alpha, \beta \in X(\Sigma)$. If $d(\alpha, \beta) \leq 2$ then there is only one fully tight multigeodesic connecting α and β and therefore only finitely many tight geodesics. When $d(\alpha, \beta) \geq 3$, by Theorem 1.4 the vertices of any tight multigeodesic, and therefore any tight geodesic, connecting α and β have intersection number with α and intersection number with β at most $F(\iota(\alpha, \beta))$. It follows from Lemma 5.1 with $K = 2F(\iota(\alpha, \beta))$ that only finitely many curves can appear on tight geodesics connecting α and β . There are thus only finitely many tight geodesics connecting α and β . This completes the proof of Theorem 1.2. QED

Remark All the dependence on the genus $g(\Sigma)$ is subsumed by $\iota(\alpha, \beta)$ since $\{\alpha, \beta\}$ is assumed to fill the surface.

6 Proof of Theorem 1.5

An algorithm is given as follows. If $d(\alpha, \beta) \leq 2$ there is only one fully tight multigeodesic connecting α and β . We assume $d(\alpha, \beta) \geq 3$ so that α and β together fill the surface. By Lemma 5.1, the set of all curves γ found on tight geodesics connecting α and β and satisfying $\iota(\alpha, \gamma) + \iota(\gamma, \beta) \leq F(\iota(\alpha, \beta))$ has finite cardinality uniformly and explicitly bounded above in terms of $\iota(\alpha, \beta)$ and $n(\Sigma)$. We list all of the paths connecting α and β of length at most $\lfloor 2\log_2(\iota(\alpha, \beta)) \rfloor + 2$ that can be constructed using these curves. By Lemma 2.3 there exists at least one such path. The shortest paths on this list are all geodesics, and those that are tight account for all of the tight geodesics. This completes the proof of Theorem 1.5. QED

7 Stable lengths

In (3) hyperbolicity constants for $\mathcal{G}(\Sigma)$ in terms of the logarithm of the surface complexity are found. Theorem 1.4 of (4) asserts the existence of a positive integer N depending only on the topology of the surface such that, for each pseudo-Anosov mapping class ϕ , there exists a tight geodesic axis in $\mathcal{G}(\Sigma)$ invariant under the action of ϕ^N .

Proposition 7.1 *There exists a finite time algorithm which takes as input a surface Σ , a pseudo-Anosov mapping class ϕ and a positive integer N such that ϕ^N has a geodesic axis and returns the stable length $\|\phi\|$ of ϕ .*

Proof We may fix a choice of k such that $\mathcal{G}(\Sigma)$ is k -hyperbolic (23), and choose an integer $M \geq 50k + 20$. If ϕ is not already the N th power of a pseudo-Anosov mapping class we may replace it with such and remember to divide the new stable length by N . Thus, without loss of generality, let us suppose that ϕ is already the N th power of some pseudo-Anosov mapping class. Then, ϕ is also a pseudo-Anosov mapping class and has a geodesic axis, denoted L , on which it acts by translation.

Let α be any curve and use Corollary 1.5 to construct a geodesic ρ from α to $\phi^M \alpha$ in $\mathcal{G}(\Sigma)$. In a k -hyperbolic geodesic metric space each geodesic rectangle is $8k$ -thin, so that any point on any one side of the rectangle is within distance $8k$ from the union of the other three; see (11; 10). A midpoint or near midpoint β of ρ lies within $8k$ of L . On the one hand we have $\|\phi\| = \|\phi^M\|/M \leq d(\beta, \phi^M \beta)/M$, and on the other

$$\begin{aligned} d(\beta, \phi^M \beta)/M &\leq (\|\phi^M\| + 16k)/M \\ &= (M\|\phi\| + 16k)/M \\ &= \|\phi\| + (16k/M) \\ &< \|\phi\| + 1. \end{aligned}$$

That is, $d(\beta, \phi^M \beta)/M - 1 < \|\phi\|$. Combining the two inequalities, we have

$$d(\beta, \phi^M \beta)/M - 1 < \|\phi\| \leq d(\beta, \phi^M \beta)/M$$

and this uniquely determines the integer $\|\phi\|$. To compute $\|\phi\|$ we now need only appeal to Corollary 1.6 and calculate $d(\beta, \phi^M \beta)$. \square

We have neither constructed an axis for ϕ^N nor have we seen how to compute appropriate values for N . It may interest the reader to find ways of doing so.

References

- [1] R. Benedetti, C. Petronio, *Lectures on Hyperbolic Geometry* : Springer Universitext (1992).
- [2] M. Bestvina, K. Fujiwara, *Bounded cohomology of subgroups of the mapping class groups* : *Geom. Topol.* **6** (2002) 69–89.
- [3] B. H. Bowditch, *Intersection numbers and the hyperbolicity of the curve complex* : *J. reine angew. Math.* **598** (2006) 105–129.
- [4] B. H. Bowditch, *Tight geodesics in the curve complex* : *Invent. Math.* **171** (2008) 281–300.
- [5] J. F. Brock, R. D. Canary, Y. N. Minsky, *The classification of Kleinian surface groups, II: The Ending Lamination Conjecture* : arXiv:math/0412006v1.
- [6] A. J. Casson, S. A. Bleiler, *Automorphisms of surfaces after Nielsen and Thurston* : London Math. Soc., Student Texts **9** (1982).
- [7] B. Farb, D. Margalit, *A primer on mapping class groups* : to appear, P. U. P. (December 2010).
- [8] A. Fathi, F. Laudenbach, V. Poénaru, *Travaux de Thurston sur les surfaces (seconde édition)* : Asterisque **66-67**, Société mathématique de France (1991).
- [9] C. D. Feustel, *Homotopic arcs are isotopic* : *Proc. Amer. Math. Soc.* **17** No. 4 (1966) 891–896.
- [10] E. Ghys, P. de la Harpe, *Sur les groupes hyperboliques d’après Mikhael Gromov* : Birkhäuser (1990).
- [11] M. Gromov, *Hyperbolic groups* : Essays in group theory (ed. S. M. Gersten), MSRI Publications **8**, Springer-Verlag (1987).
- [12] J. L. Harer, *The virtual cohomological dimension of the mapping class groups of an orientable surface* : *Invent. Math.* **84** (1986) 157–176.
- [13] 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. Math. Stud.* No. 97, P. U. P. (1981) 245-251.
- [14] J. P. Hempel, *3-manifolds as viewed from the curve complex* : *Topology* **40** (2001) 631–657.
- [15] N. V. Ivanov, *Automorphisms of complexes of curves and of Teichmüller spaces* : *I. M. R. N.* (1997) 651–666.
- [16] J. Jost, *Compact Riemann surfaces* : Springer Universitext, third edition (2006).

- [17] R. P. Kent IV : private communication (December 2004).
- [18] M. Korkmaz, *Automorphisms of complexes of curves on punctured spheres and on punctured tori* : Topology App. **95** (1999) 85–111.
- [19] J. P. Leasure, *Geodesics in the complex of curves of a surface* : PhD thesis University of Texas (2002) available on the Texas at Austin Digital Archive from January 2006, URL <http://hdl.handle.net/2152/1700>.
- [20] H. I. Levine, *Homotopic curves on surfaces* : Proc. Amer. Math. Soc. **14** No. 6 (1963) 986–990.
- [21] W. B. R. Lickorish, *A representation of orientable combinatorial 3-manifolds* : Ann. Math. (2) **76** No. 3 (1962) 531–540.
- [22] F. Luo, *Automorphisms of the complex of curves* : Topology **39** (2000) 283–298.
- [23] H. A. Masur, Y. N. Minsky, *Geometry of the complex of curves I: Hyperbolicity* : Invent. Math. **138** (1999) 103–149.
- [24] H. A. Masur, Y. N. Minsky, *Geometry of the complex of curves II: Hierarchical structure* : Geom. Funct. Anal. **10** (2000) 902–974.
- [25] Y. N. Minsky, *The classification of Kleinian surface groups, I: Models and bounds* : Ann. Math. **1** (2010) 1–107.
- [26] I. Richards, *On the classification of non-compact surfaces* : Trans. Amer. Math. **106** (1963) 259–269.
- [27] H. Seifert, W. Threlfall, *Lehrbuch der Topologie* : Teubner, Leipzig (1934).
- [28] K. J. Shackleton, *An acylindricity theorem for the mapping class group* : New York J. Math. **16** (2010) 563–573.

University of Tokyo IPMU, 5-1-5 Kashiwa-no-ha, Kashiwa-shi, Chiba, 277-8568 Japan

e-mail: kenneth.shackleton@ipmu.jp

e-mail: kjs2006@alumni.soton.ac.uk

<http://member.ipmu.jp/kenneth.shackleton/>

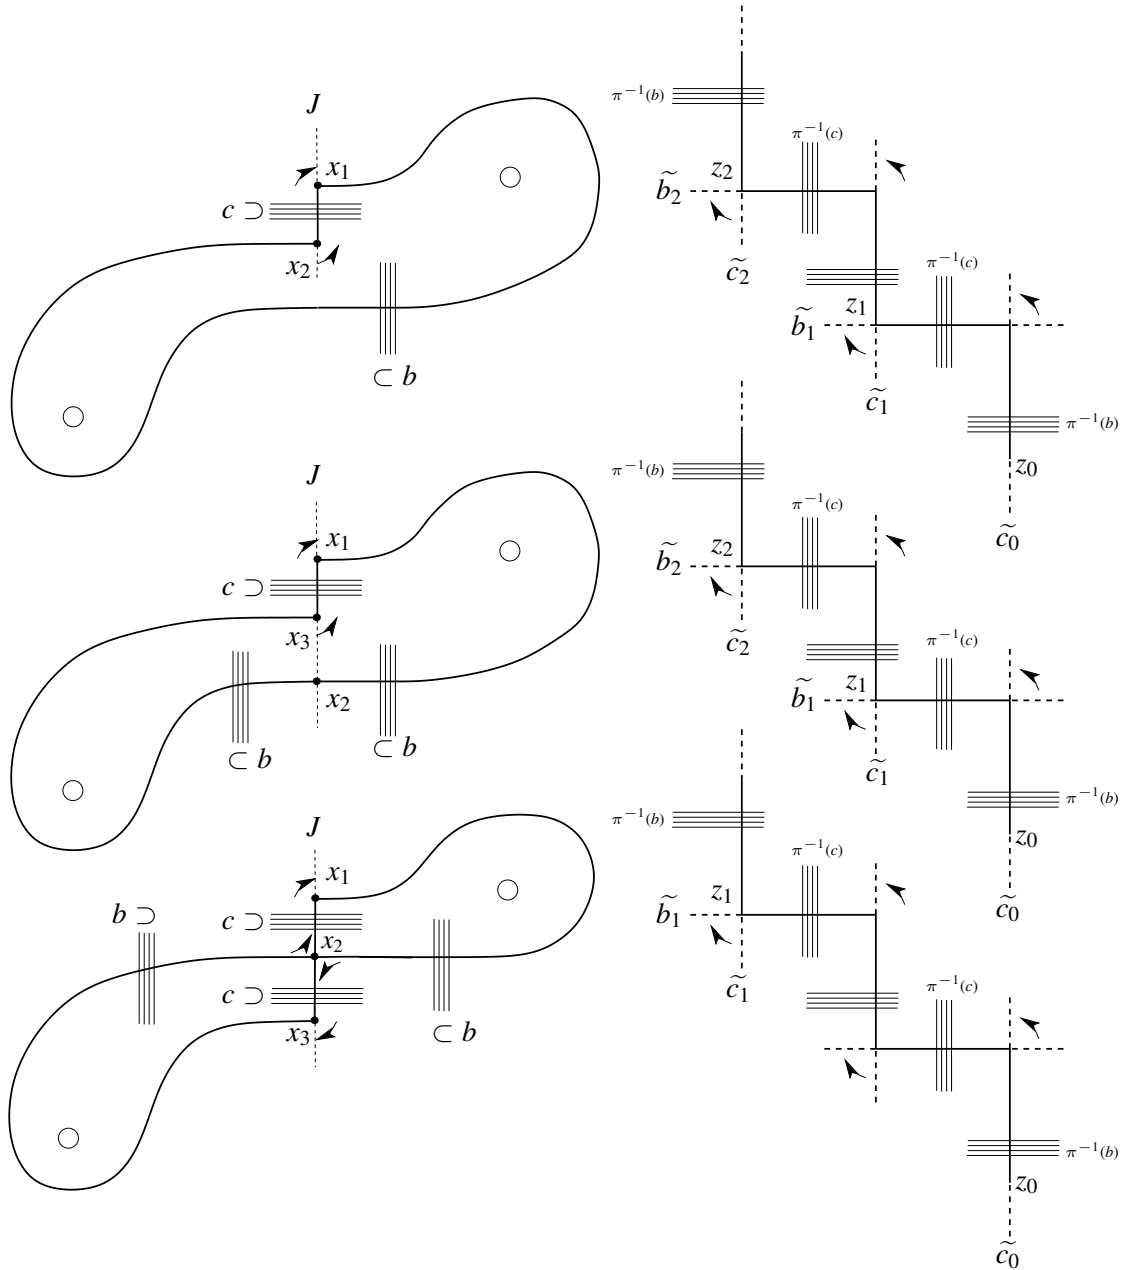


Figure 1: The three types of loop, from Type 1 down to Type 3, defining p (left) and compact pieces of its lifts (right) marching “left-right” in the hyperbolic plane.




Figure 2: The subarc h of v_m is necessarily disjoint from the multiloop v_0 .