Math. Res. Lett. **15** (2008), *no.* 2, 309–320

TOTALLY GEODESIC SUBGRAPHS OF THE PANTS COMPLEX

Javier Aramayona, Hugo Parlier, and Kenneth J. Shackleton

Abstract. In recent work of Brock's, the pants graph is shown to be a combinatorial model for the completion of the Weil-Petersson metric on Teichmüller space. We prove that every Farey graph embedded in the pants graph is totally geodesic, in analogy with the extrinsic geometry of any 2-dimensional stratum inside the Weil-Petersson completion.

1. Introduction

Let Σ be a compact, connected and orientable surface, possibly with non-empty boundary, of genus $g(\Sigma)$ and $|\partial\Sigma|$ boundary components, and refer to as the *mapping class group* $\mathrm{Map}(\Sigma)$ the group of all self-homeomorphisms of Σ up to homotopy.

After Hatcher-Thurston [5], to the surface Σ one may associate a simplicial graph $\mathcal{P}(\Sigma)$, the *pants graph*, whose vertices are all the pants decompositions of Σ and any two vertices are connected by an edge if and only if they differ by an elementary move; see Section 2.2 for an expanded definition. This graph is connected, and one may define a path-metric d on $\mathcal{P}(\Sigma)$ by first assigning length 1 to each edge and then regarding the result as a length space.

The pants graph, with its own geometry, is a fundamental object to study. Brock [2] revealed deep connections with hyperbolic 3-manifolds and proved the pants graph is the correct combinatorial model for the Weil-Petersson metric on Teichmüller space, for the two are quasi-isometric. The isometry group of $(\mathcal{P}, d)$ is also correct in so far as the study of surface groups is concerned, for Margalit [6] proved it is almost always isomorphic to the mapping class group of Σ. In addition, Masur-Schleimer [9] proved the pants graph to be one-ended for closed surfaces of genus at least 3. With only a few exceptions, the pants graph is not hyperbolic in the sense of Gromov [3].

Our main result concerns the geometry of the pants graph.

Theorem 1. *Let Σ be a compact, connected and orientable surface, and denote by $\mathcal{F}$ a Farey graph. Let $\phi\colon \mathcal{F} \to \mathcal{P}(\Sigma)$ be a simplicial embedding. Then, $\phi(\mathcal{F})$ is totally geodesic in $\mathcal{P}(\Sigma)$.*

The completion of the Weil-Petersson metric can be characterised by attaching so-called strata [7]. These are totally geodesic subspaces of the completion, by a result of Wolpert [10], and correspond to lower dimensional Teichmüller spaces, or products thereof, each with their own Weil-Petersson metric, or product metric. Combining this with Theorem 1.1 of Brock [2], one finds the Farey subgraphs of the pants graph are uniformly quasi-isometrically embedded. This fact is also implicit in the earlier work of Masur-Minsky [8].

Received by the editors February 26, 2007.
Key words and phrases. pants complex; Weil-Petersson metric; Farey graph.

Theorem 1 offers a complete analogy between the geometry of the Farey subgraphs in a pants graph and the geometry of the corresponding strata lying in the completed Weil-Petersson space. In order to prove Theorem 1, we shall use Theorem 2 to project paths in the pants graph to paths in the given Farey graph of no greater length. All the notation of Theorem 2 will be explained in Section 2, but for now we point out the finite set of curves $\pi_Q(\nu)$ is the subsurface projection after Masur-Minsky [8] of a pants decomposition ν to the Farey graph $\mathcal{F}_Q$ determined by the codimension 1 multicurve Q. The intrinsic metric on this Farey graph, assigning length 1 to each edge, is denoted by d_Q.

Theorem 2. *Let Σ be a compact, connected and orientable surface and denote by Q a codimension 1 multicurve on Σ. Let $(\nu_0, \dots, \nu_n)$ be a path in the pants graph $\mathcal{P}(\Sigma)$. For each index $i \leq n-2$ and for each $\delta_i \in \pi_Q(\nu_i)$, there exists an integer $j \in \{1, 2\}$ and a curve $\delta_{i+j} \in \pi_Q(\nu_{i+j})$ such that $d_Q(\delta_i, \delta_{i+j}) \leq j$.*

To the authors' knowledge, it has yet to be decided whether there exists a distance non-increasing projection from the whole pants graph to any one of its Farey subgraphs. In the absence of an affirmative result, Theorem 2 may well hold independent interest.

Let us indicate two consequences of Theorem 1. First, note that for any hyperbolic self-isometry f of a Farey graph, there exists a bi-infinite geodesic invariant under the action of f^2.

Corollary 3. *Let $f \in \mathrm{Map}(\Sigma)$ be any mapping class leaving invariant a subgraph of $\mathcal{P}(\Sigma)$ isomorphic to a Farey graph, on which it acts as a hyperbolic self-isometry. Then, there exists a bi-infinite geodesic in $\mathcal{P}(\Sigma)$ invariant under the action of f^2.*

We remark that examples of such mapping classes include those whose restriction to the complement of some complexity 1 subsurface Y is the identity and whose restriction to Y is a pseudo-Anosov mapping class.

Second, let Q be a multicurve on Σ with the property that every complementary component of Q is a surface of complexity 1. Then, the subgraph of $\mathcal{P}(\Sigma)$ spanned by all pants decompositions containing Q is isomorphic to a product of Farey graphs, each totally geodesic by Theorem 1. By considering one bi-infinite geodesic in each Farey graph, we deduce the following. Note, by a *line* in the free abelian group $\mathbb{Z}^r$ we shall mean a coset of any one of the $\mathbb{Z}$-factors.

Corollary 4. *Let r denote the integer $\lfloor (3g(\Sigma) + |\partial\Sigma| - 2)/2 \rfloor$. There exists a quasi-isometric embedding from $\mathbb{Z}^r$, given the L^1-metric, into $\mathcal{P}(\Sigma)$ such that the image of any line is a geodesic.*

Thus, infinitely many of the maximal quasi-flats in $\mathcal{P}(\Sigma)$ identified by the Geometric Rank Theorem [3, 1, 4] are convex in their principal directions. However, establishing the existence of convex maximal flats remains an open problem.

The plan of this paper is as follows. In Section 2 we recall all the terminology we need, much of which is already standard. In Section 3 we give an elementary proof to Theorem 2. Indeed, if Q borders a 1-holed torus or a 4-holed sphere with only one essential boundary component, it transpires that one may always take $j = 1$. In Section 4 we apply Theorem 2 to give an elementary proof to Theorem 1.

Let us close the introduction by stating the following conjecture.

Conjecture 5. *Let Σ_1 and Σ_2 be a pair of compact and orientable surfaces. Let $\phi\colon \mathcal{P}(\Sigma_1) \to \mathcal{P}(\Sigma_2)$ be a simplicial embedding. Then, $\phi(\mathcal{P}(\Sigma_1))$ is totally geodesic in $\mathcal{P}(\Sigma_2)$.*

2. Background and definitions

We supply all the background and terminology needed both to understand the statements of our main results, and to make sense of their proofs. Throughout, we define a *loop* on Σ as the homeomorphic image of a standard circle.

2.1. Curves and multicurves. A loop on Σ is said to be *trivial* only if it bounds a disc and *peripheral* only if it bounds an annulus whose other boundary component belongs to $\partial\Sigma$. For a non-trivial and non-peripheral loop c, we shall denote by $[c]$ its free homotopy class. A *curve* is by definition the free homotopy class of a non-trivial and non-peripheral loop. Given any two curves α and β, their *intersection number* $\iota(\alpha, \beta)$ is defined equal to $\min\{|\mathrm{a} \cap \mathrm{b}| : \mathrm{a} \in \alpha, \mathrm{b} \in \beta\}$.

We shall say two curves are *disjoint* only if they have zero intersection number, and otherwise say they *intersect essentially.* A pair of curves $\{\alpha, \beta\}$ is said to *fill* the surface Σ only if $\iota(\delta, \alpha) + \iota(\delta, \beta) > 0$ for every curve δ. In other words, every curve on Σ intersects at least one of α and β essentially.

A *multicurve* is a collection of distinct and disjoint curves, and the intersection number for a pair of multicurves is to be defined additively. We denote by $\kappa(\Sigma)$ the cardinality of any maximal multicurve on Σ, equal to $3g(\Sigma) + |\partial\Sigma| - 3$, and refer to this as the *complexity of* Σ. Note, the only surfaces of complexity 1 are the 4-holed sphere and the 1-holed torus.

Given a set of disjoint loops L, such as the boundary of some subsurface of Σ, we denote by $[L]$ the multicurve maximal among all multicurves whose every curve is represented by some element of L. We shall say a multicurve ω has *codimension* k, for some non-negative integer k, only if $|\omega| = \kappa(\Sigma) - k$.

2.2. Pants decompositions. A *pants decomposition* of a surface is a maximal collection of distinct and disjoint curves, in other words a maximal multicurve. Two pants decompositions μ and ν are said to be related by an *elementary move* only if $\mu \cap \nu$ is a codimension 1 multicurve and the remaining two curves together either fill a 4-holed sphere and intersect twice or fill a 1-holed torus and intersect once; consider Figure 1.

2.3. Arcs. An *arc* on Σ is the homotopy class, relative to $\partial\Sigma$, of an embedded interval ending on $\partial\Sigma$ that does not bound a disc with $\partial\Sigma$. There are broadly two types of arc: those that end on only one component of $\partial\Sigma$, referred to as *waves*, and those that end on two different components of $\partial\Sigma$, referred to as *seams*; see Figure 2 below.

Typically, our arcs will live on proper subsurfaces of complexity 1, noting every arc on a 1-holed torus is a wave. We may similarly define the intersection number of a pair of arcs, or an arc and a curve, and say two arcs are disjoint or intersect essentially.

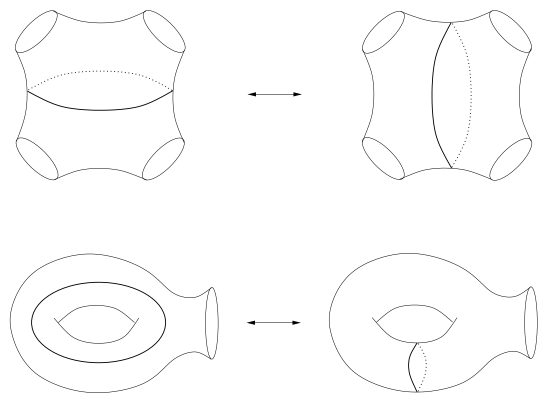

Figure 1. The two types of elementary move.

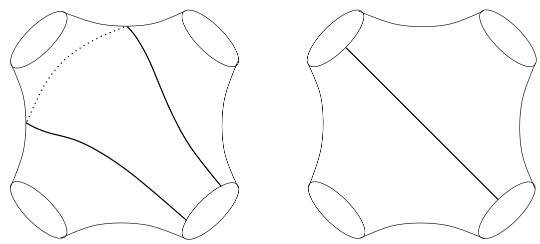

Figure 2. The two types of arc, respectively a wave and a seam.

2.4. Graphs and paths. For us, a *path* in a graph shall be a finite sequence of vertices such that any consecutive pair spans an edge; one can recover a topological path by joining up the dots. A *geodesic* is then a path realising distance. Finally, a subgraph F of a graph G is said to be *totally geodesic* only if every geodesic in G whose two endpoints belong to F is in fact entirely contained in F.

2.5. Farey graphs. There are numerous ways to build a Farey graph $\mathcal{F}$, any two producing isomorphic graphs. One can start with the rational projective line $\widehat{\mathbb{Q}} := \mathbb{Q} \cup \{\infty\}$, identifying 0 with $\frac{0}{1}$ and ∞ with $\frac{1}{0}$, and take this to be the vertex set of $\mathcal{F}$. Then, two projective rational numbers $\frac{p}{q}, \frac{r}{s} \in \widehat{\mathbb{Q}}$, where p and q are coprime and r and s are coprime, are deemed to span an edge, or 1-simplex, if and only if $|ps - rq| = 1$. The result is a connected graph in which every edge separates. The graph $\mathcal{F}$ can be represented on a disc; see Figure 3 below. We shall say a graph *is a Farey graph* if it is isomorphic to $\mathcal{F}$.

It should be noted that both the pants graph of the 4-holed sphere and the pants graph of the 1-holed torus are Farey graphs. It follows that any codimension 1 multicurve Q on Σ determines a unique Farey graph $\mathcal{F}_Q$ in $\mathcal{P}(\Sigma)$; the converse is Lemma 6 from Section 3. We shall always denote by d_Q the intrinsic combinatorial metric on $\mathcal{F}_Q$, assigning length 1 to each edge.

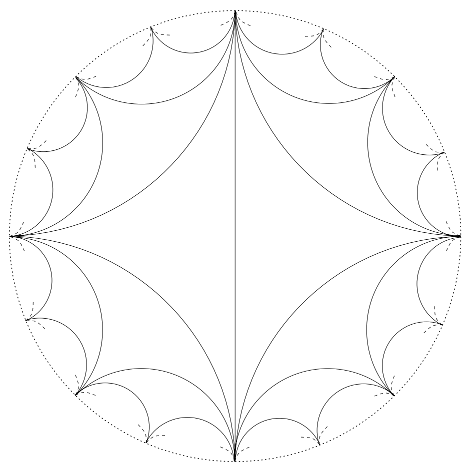

FIGURE 3. The Farey graph can be represented on a disc.

2.6. Subsurface projections. Given a curve α and an incompressible subsurface Y of Σ, we shall write $\alpha \subset Y$ only if α can be represented by a non-peripheral loop on Y. If every loop representing α has non-empty intersection with Y we can say *α and Y intersect*, otherwise we say they are *disjoint*. If every loop representing α intersects Y in at least one interval, we can say *α crosses Y*. We may similarly speak of a multicurve crossing Y, if one of its curves crosses Y.

For a codimension 1 multicurve Q, let Y denote the unique, up to isotopy, complexity 1 incompressible subsurface of Σ such that each curve in Q is disjoint from Y. Let α be any curve intersecting Y, and choose any simple representative $c \in \alpha$ such that $\#(c \cap \partial Y)$ is minimal. We refer to each component of $c \cap Y$ as a *footprint of c on Y*, and to the homotopy class of such a footprint as a *footprint of α on Y*. Note, footprints of a curve can be arcs or curves.

Given a footprint b for the curve α there exists a unique curve on Y disjoint from b, and such a curve shall be referred to as a *projection of α*. Note the set of all α projections, each counted once, depends only on α and the original multicurve Q, and we denote this set by $\pi_Q(\alpha)$. For a multicurve ν we define $\pi_Q(\nu)$ to be equal to the union $\bigcup_\nu \pi_Q(\alpha)$. The set $\pi_Q(\nu)$ is an example of a *subsurface projection*, as defined by Masur-Minsky in Section 1.1 of [8]. See Figure 4 below for an illustration.

Remark. *We note that $\pi_Q(Q) = \emptyset$ and that $\pi_Q(\delta) = \{\delta\}$ for any curve $\delta \subset Y$. Moreover, if $\delta \subset Y$ is a curve and ν is a multicurve crossing Y and disjoint from δ, then $\pi_Q(\nu) = \{\delta\}$.*

3. Proof of Theorem 2

Let us start with two elementary results, the first characterising the Farey subgraphs of any given pants graph and the second relating low intersection numbers to distances for a pair of curves on a 4-holed sphere.

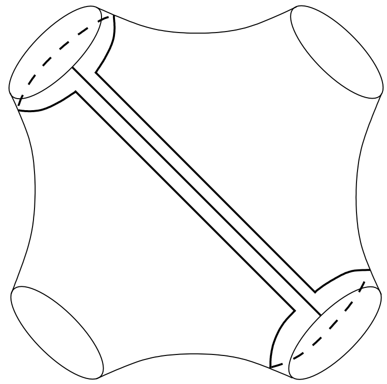

FIGURE 4. A seam and a projected curve.

Lemma 6. *Let $\phi\colon \mathcal{F} \to \mathcal{P}(\Sigma)$ be a simplicial embedding. Then, there exists a codimension 1 multicurve on Σ contained in every vertex of $\phi(\mathcal{F})$.*

Proof. This is a consequence of the following two remarks. First, note the vertices of any 3-cycle from $\phi(\mathcal{F})$ always intersect in a common codimension 1 multicurve. Second, note for any two vertices μ and ν of $\phi(\mathcal{F})$ there exists a finite sequence of 3-cycles $\Delta_0, \ldots, \Delta_n$ such that μ is a vertex of Δ_0, such that ν is a vertex of Δ_n, and such that $\Delta_i \cap \Delta_{i+1}$ is an edge for each index i. One can then prove Lemma 6 by an induction. □

Lemma 7. *Let Y be a 4-holed sphere. Then, any two vertices δ_0, δ_2 of $\mathcal{P}(Y)$ of intersection number at most 4 are at distance $d(\delta_0, \delta_2)$ at most 2.*

Proof. There exists a curve δ_1 on Y such that $\iota(\delta_1, \delta_j) \leq 2$ for both $j \in \{0, 2\}$; if $\iota(\delta_0, \delta_2) = 4$ then such a curve can be explicitly constructed by performing a standard surgery on either of δ_0 or δ_2. It follows that $d(\delta_0, \delta_2) \leq d(\delta_0, \delta_1) + d(\delta_1, \delta_2) \leq 1 + 1 = 2$. □

The following two results shall be applied in what will become known as Case B for the 4-holed sphere, Lemma 9 playing an especially important role.

Lemma 8. *Let P be a pants decomposition of Σ, and let Y be a connected complexity 1 incompressible subsurface of Σ. If P does not contain $[\partial Y]$, then P contains at least two curves intersecting Y.*

Proof. We shall denote by $\kappa_*(Y)$ the cardinality of a maximal multicurve on Σ whose every curve does not intersect Y. Let $\omega \subset P$ be the set of all curves in P that do not intersect Y. We have

$$|P| = \kappa(\Sigma) = \kappa(Y) + \kappa_*(Y) = 1 + \kappa_*(Y) \geq 1 + |\omega| + 1 = 2 + |\omega|,$$

and so $|P| \geq 2 + |\omega|$ as required. □

Lemma 9. *Let P be a pants decomposition of Σ, and let Y be an incompressible subsurface of Σ homeomorphic to a 4-holed sphere. If there exist two distinct curves in $[\partial Y]$ not contained in P, then P contains at least three curves intersecting Y.*

Proof. Let $\omega \subset P$ be the set of all curves in P that do not intersect Y. We have

$$|P| = \kappa(\Sigma) = \kappa(Y) + \kappa_*(Y) = 1 + \kappa_*(Y) \geq 1 + |\omega| + 2 = 3 + |\omega|,$$

and so $|P| \geq 3 + |\omega|$ as required. □

We now turn to proving Theorem 2, denoting by Y the complexity 1 subsurface of Σ complementary to Q. Note the statement of Theorem 2 holds vacuously if $\kappa(\Sigma) \leq 0$ and trivially if $\kappa(\Sigma) = 1$, since then ϕ is an isomorphism. When $\kappa(\Sigma) = 2$, the surface Σ is either a 5-holed sphere or a 2-holed torus. If the genera $g(Y)$ and $g(\Sigma)$ are equal, then each footprint of ν_{i+1} on Y is therefore either a curve or a wave. As such, there exists a curve $\delta_{i+1} \in \pi_Q(\nu_{i+1})$ such that δ_i and δ_{i+1} are either equal or intersect minimally. We can then take $j = 1$, noting $d_Q(\delta_i, \delta_{i+1}) = 1$. The remaining case, Σ the 2-holed torus and Y the 4-holed sphere, is deferred to Appendix.

For the remainder of this section, it is to be assumed that $\kappa(\Sigma) \geq 3$. Let $\delta_i \in \pi_Q(\nu_i)$. In constructing a curve δ_{i+1} or δ_{i+2}, as per the statement of Theorem 2, we note Lemma 6 tells us it is enough to consider separately the case Y is a 4-holed sphere and the case Y is a 1-holed torus.

Y IS A 4-HOLED SPHERE

The case Y is a 4-holed sphere separates into two main cases, according as δ_i belongs to ν_i or does not belong to ν_i.

Case A: $\delta_i \in \nu_i$.

I. $\delta_i \in \nu_{i+1}$. Take $j = 1$ and $\delta_{i+1} = \delta_i$.

II. $\delta_i \notin \nu_{i+1}$. We may still take $j = 1$ and choose any $\delta_{i+1} \in \pi_Q(\nu_{i+1})$, for δ_i is a curve and, as such, is disjoint from $[\partial Y]$. Now $\iota(\delta_i, \delta_{i+1}) \leq 2$ and so $d_Q(\delta_i, \delta_{i+1}) \leq 1$.

Case B: $\delta_i \notin \nu_i$.

By definition, there exists a Y-footprint a_i of ν_i such that $\iota(a_i, \delta_i) = 0$. We denote by α_i any curve from ν_i having a_i as a footprint. Let a_{i+1} be any footprint of ν_{i+1} on Y, and let α_{i+1} be any element of ν_{i+1} having a_{i+1} as a footprint.

I. *a_i and a_{i+1} intersect essentially.* Since a_i and a_{i+1} intersect essentially, so must the two curves α_i and α_{i+1}. Moreover, as $\delta_i \notin \nu_i$, so a_i can only be an arc.

Suppose first that a_{i+1} is a curve. Then, α_{i+1} and a_{i+1} are equal. According to Lemma 8 there exists a curve $\alpha'_i \in \nu_i$ such that $\alpha_i \neq \alpha'_i$ and such that α'_i intersects Y. Since $d(\nu_i, \nu_{i+1}) = 1$ and since $\iota(\alpha_i, \alpha_{i+1}) \neq 0$, so $\alpha'_i \in \nu_{i+1}$. The set $\{\alpha'_i, \alpha_{i+1}\} \cap \nu_{i+2}$ is therefore non-empty. Let $\gamma \in \{\alpha'_i, \alpha_{i+1}\} \cap \nu_{i+2}$ and take $j = 2$. There exists $\delta_{i+2} \in \pi_Q(\gamma)$ such that $\iota(\delta_i, \delta_{i+2}) \leq 4$ and so, according to Lemma 7, $d_Q(\delta_i, \delta_{i+2}) \leq 2$.

Henceforth, a_{i+1} shall always be an arc. Appealing to Lemma 8, there exists a Y-footprint a'_{i+1} of ν_{i+1} and a corresponding curve $\alpha'_{i+1} \in \nu_{i+1}$ such that $\alpha'_{i+1} \neq \alpha_{i+1}$. Since $d(\nu_i, \nu_{i+1}) = 1$ it follows that $\iota(a_i, a'_{i+1}) = 0$. Note, if a'_{i+1} is a curve then $\alpha'_{i+1} = a'_{i+1}$ and we may take $j = 1$ and $\delta_{i+1} = \alpha'_{i+1}$.

Henceforth, a'_{i+1} is assumed to be an arc. We observe a_{i+1} and a'_{i+1} are distinct arcs, since a_{i+1} intersects a_i essentially whereas a'_{i+1} is disjoint from a_i. The first case, B.I., will now be completed by considering in turn the two topological possibilities for a_i.

I.(i) *a_i is a wave.* Let $\gamma'_{i+1} \in \pi_Q(\nu_{i+1})$ be such that $\iota(\gamma'_{i+1}, a'_{i+1}) = 0$. Then, $\iota(\delta_i, \gamma'_{i+1}) \leq 2$ and therefore $d_Q(\delta_i, \gamma'_{i+1}) \leq 1$. We can thus take $j = 1$ and $\delta_{i+1} = \gamma'_{i+1}$.

I.(ii) *a_i is a seam.* If in addition a'_{i+1} is a wave, then we may argue as per Case B.I(i) where the types of a_i and a'_{i+1} are interchanged. Henceforth, a'_{i+1} shall be a seam.

Suppose $\{a_i, a'_{i+1}\}$ ends on at least three different components of ∂Y. Let $\gamma'_{i+1} \in \pi_Q(\nu_{i+1})$ be such that $\iota(\gamma'_i, a'_{i+1}) = 0$. Then, $\iota(\delta_i, \gamma'_{i+1}) \leq 2$ and so $d_Q(\delta_i, \gamma'_{i+1}) \leq 1$. We now take $j = 1$ and $\delta_{i+1} = \gamma'_{i+1}$.

Suppose instead $\{a_i, a'_{i+1}\}$ now ends on at most two, and therefore exactly two, different components of ∂Y. Since a'_{i+1} is a seam and since $\alpha_{i+1} \in \nu_{i+1}$, so ν_{i+1} fails to contain at least two curves from $[\partial Y]$. If the two components of ∂Y on which a_i ends are not homotopic on Σ, then by Lemma 9 there exists a curve $\alpha''_{i+1} \in \nu_{i+1}$ such that $\alpha''_{i+1} \notin \{\alpha_{i+1}, \alpha'_{i+1}\}$ and such that α''_{i+1} intersects Y. (The remaining case, namely the two components of ∂Y on which a_i ends are homotopic, seems to require special consideration, and so we prefer to postpone this to Appendix.) Since $d(\nu_i, \nu_{i+1}) = 1$ and since $\iota(\alpha_i, \alpha_{i+1}) \neq 0$, so $\iota(\alpha_i, \alpha''_{i+1}) = 0$. Moreover, since $\alpha_i \notin \nu_{i+1}$, so $\alpha_i \neq \alpha''_{i+1}$. As $d(\nu_{i+1}, \nu_{i+2}) = 1$, so $\{\alpha'_{i+1}, \alpha''_{i+1}\} \cap \nu_{i+2} \neq \emptyset$. Let $\gamma \in \{\alpha'_{i+1}, \alpha''_{i+1}\} \cap \nu_{i+2}$. We now take $j = 2$ and $\delta_{i+2} \in \pi_Q(\gamma)$, noting that $\iota(\delta_i, \delta_{i+2}) \leq 4$ and, as such, $d_Q(\delta_i, \delta_{i+2}) \leq 2$.

II. *a_i and a_{i+1} are disjoint.* First note that, if either of a_i and a_{i+1} is a wave, we may take $j = 1$ and $\delta_{i+1} \in \pi_Q(\alpha_{i+1})$ such that $\iota(\delta_{i+1}, a_{i+1}) = 0$. Then, $\iota(\delta_i, \delta_{i+1}) \leq 2$ and, as such, $d_Q(\delta_i, \delta_{i+1}) \leq 1$. Henceforth, we assume that a_i and a_{i+1} are both seams.

If $\{a_i, a_{i+1}\}$ ends on at least three components of ∂Y we may take $j = 1$ and $\delta_{i+1} \in \pi_Q(\alpha_{i+1})$ such that $\iota(\delta_{i+1}, a_{i+1}) = 0$. Then, $\iota(\delta_i, \delta_{i+1}) \leq 2$ and, as such, $d_Q(\delta_i, \delta_{i+1}) \leq 1$.

Thus, we may assume that $\{a_i, a_{i+1}\}$ ends on at most two, and therefore exactly two, components of ∂Y. By assumption, ν_{i+1} does not contain $[\partial Y]$. According to Lemma 8, there exists a second Y-footprint a'_{i+1} for some curve $\alpha'_{i+1} \in \nu_{i+1}$ such that α_{i+1} and α'_{i+1} are distinct. If a_i and a'_{i+1} are equal then $\delta_i \in \pi_Q(\nu_{i+1})$, and we may take $j = 1$ and $\delta_{i+1} = \delta_i$. We may therefore assume a_i and a'_{i+1} are not equal.

If a_i and a'_{i+1} intersect essentially, then we may appeal to Case B.I. with a'_{i+1} substituted for a_{i+1}. We may thus assume that a_i and a'_{i+1} are disjoint. Since three homotopically distinct and disjoint arcs on Y cannot end on at most two components of ∂Y, it follows $\{a_i, a'_{i+1}\}$ ends on at least three different components of ∂Y. We can now take $j = 1$ and $\delta_{i+1} \in \pi_Q(\alpha'_{i+1})$ such that $\iota(\delta_{i+1}, a'_{i+1}) = 0$.

This concludes the case Y is a 4-holed sphere.

Y IS A 1-HOLED TORUS

The case of the 1-holed torus is more straightforward, for here each arc is a wave, and can be treated by considering separately four mutually exclusive cases.

I. *ν_i, ν_{i+1} contain $[\partial Y]$.* Let δ_{i+1} denote the only curve contained in $\pi_Q(\nu_{i+1})$. We may then take $j = 1$ and note $d_Q(\delta_i, \delta_{i+1})) \leq 1$.

II. *ν_i contains $[\partial Y]$, whereas ν_{i+1} does not.* Then, $\delta_i \in \nu_{i+1}$. We may take $j = 1$ and $\delta_{i+1} = \delta_i$.

III. *ν_{i+1} contains $[\partial Y]$, whereas ν_i does not.* Then, ν_{i+1} contains a single curve γ_{i+1} such that $\gamma_{i+1} \subset Y$. Since $d(\nu_i, \nu_{i+1}) = 1$, so $\gamma_{i+1} \in \nu_i$ and hence $\gamma_{i+1} \in \pi_Q(\nu_i)$. As $\pi_Q(\nu_i)$ contains only one element, so $\gamma_{i+1} = \delta_i$. We may now take $j = 1$ and $\delta_{i+1} = \delta_i$.

IV. *Neither ν_i nor ν_{i+1} contains $[\partial Y]$.* By definition, there exists a Y-footprint a_i of ν_i such that $\iota(\delta_i, a_i) = 0$. According to Lemma 8, there exist two footprints a_{i+1} and a'_{i+1} of ν_{i+1} corresponding to different elements of ν_{i+1}. Since $d(\nu_i, \nu_{i+1}) = 1$, so at least one of these footprints, say a_{i+1}, is disjoint from a_i. We may take $j = 1$ and $\delta_{i+1} \in \pi_Q(\nu_{i+1})$ such that $\iota(\delta_{i+1}, a_{i+1}) = 0$. Note, $\iota(\delta_i, \delta_{i+1}) \leq 1$ and, as such, $d_Q(\delta_i, \delta_{i+1}) \leq 1$.

This concludes the case Y is a 1-holed torus, and a proof of Theorem 2.

4. Proof of Theorem 1

Let $\mathcal{F}$ be a Farey graph and let $\phi\colon \mathcal{F} \to \mathcal{P}(\Sigma)$ be a simplicial embedding. There exists a unique codimension 1 multicurve Q on Σ such that Q is contained in every vertex of $\phi(\mathcal{F})$; see Lemma 6.

Suppose, for contradiction, that $\phi(\mathcal{F})$ is not totally geodesic. Then, there exists a geodesic $\nu_0, \nu_1, \ldots, \nu_n$ in $\mathcal{P}(\Sigma)$ such that $\{\nu_0, \nu_n\} \subset \phi(\mathcal{F})$ but $\nu_j \notin \phi(\mathcal{F})$ for each $j \in \{1, 2, \ldots, n-1\}$. Applying Theorem 2 inductively, we find an increasing sequence of integers $\{k_1, k_2, \ldots, k_m\} \subseteq \{1, 2, \ldots, n\}$, containing 1 and at least one of $n-1$ and n, and a corresponding sequence of curves $\delta_{k_j} \in \pi_Q(\nu_{k_j})$ such that $0 < k_{j+1} - k_j \leq 2$, for each j, and such that $d_Q(\delta_{k_j}, \delta_{k_{j+1}}) \leq k_{j+1} - k_j$, for each j. Necessarily, $\phi(\delta_{k_1}) = \nu_0$ and $\phi(\delta_{k_m}) = \nu_n$, by the closing remark of Section 2.6. We note that

$$d_Q(\delta_{k_1}, \delta_{k_m}) \leq \sum_j d_Q(\delta_{k_j}, \delta_{k_{j+1}}) \leq \sum_j k_{j+1} - k_j \leq n - 1,$$

and, since paths in $\mathcal{F}$ determine paths in $\mathcal{P}(\Sigma)$ via ϕ, so it follows that

$$d(\nu_0, \nu_n) \leq d_Q(\delta_{k_1}, \delta_{k_m}) \leq n - 1.$$

This is a contradiction, and the statement of Theorem 1 follows.

Appendix

We treat separately one instance of Case B.I.(ii) from the proof of Theorem 2, where the seam a_i ends on two distinct but homotopic components of ∂Y. This simultaneously treats the case Σ is a 2-holed torus and Y is a 4-holed sphere. In either instance, we cannot appeal to Lemma 9.

We recall a_{i+1} is a footprint of $\alpha_{i+1} \in \nu_{i+1}$ on Y that intersects a_i essentially, and that a'_{i+1} is a footprint of $\alpha'_{i+1} \in \nu_{i+1}$ on Y both disjoint from and non-homotopic to a_i. In addition, we might as well assume $\{a_i, a_{i+1}\}$ and $\{a_i, a'_{i+1}\}$ both end on precisely two distinct components of ∂Y, for we may otherwise take $j = 1$ and readily find $\delta_{i+1} \in \pi_Q(\nu_{i+1})$ as claimed.

I. a_{i+1} *is a seam.* Only one subcase, up to symmetry and depicted in the upper-left diagram of Figure 5, is legal. For this subcase alone, let $\gamma \in \pi_Q(\alpha_{i+1})$ be the curve such that $\iota(\gamma, a_{i+1}) = 0$. Then, $\iota(\delta_i, \gamma) = 8$. However, there exists a further curve $\gamma' \subset Y$ such that $\iota(\delta_i, \gamma') = 2$ and $\iota(\gamma', \gamma) = 2$. Thus $(\delta_i, \gamma', \gamma)$ is a path in $\mathcal{F}_Q$, and it follows $d_Q(\delta_i, \gamma) \leq 2$, in fact precisely 2. We may therefore take $j = 2$ and find $\delta_{i+2} \in \pi_Q(\{\alpha_{i+1}, \alpha'_{i+1}\} \cap \nu_{i+2})$ such that $d_Q(\delta_i, \delta_{i+2}) \leq 2$.

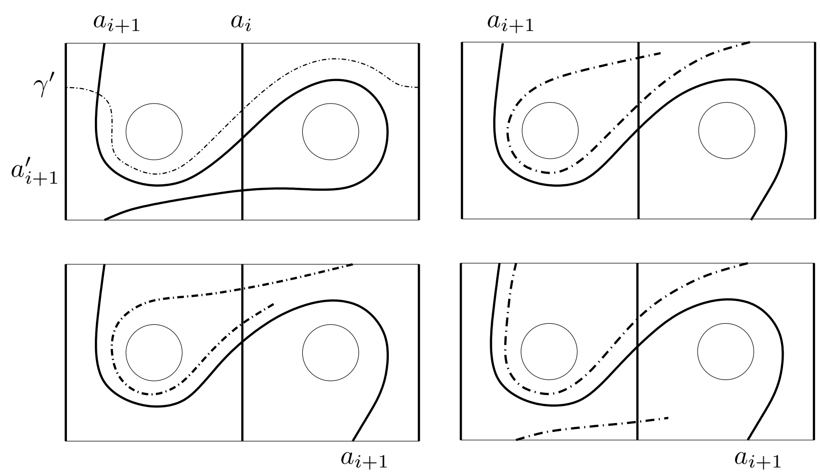

FIGURE 5. The four subcases arising when a_{i+1} is a seam, up to symmetry, of which only that depicted in the upper-left diagram is legal. In each diagram we represent the 4-holed sphere Y as a disc with four corners and two holes, identifying the left and right vertical edges to give a'_{i+1} and the middle vertical edge with a_i. The top and bottom edges correspond to distinct components of ∂Y, homotopic on Σ. The central arc a_{i+1} cuts the top and bottom edges in two; we have normalised so that the arc a_{i+1} is always incident on the upper-left half-edge.

In each of the remaining subcases, any attempt to complete the arc a_{i+1} to the curve α_{i+1}, such as that indicated by a broken line, either fails to produce a simple closed curve or produces a curve of intersection number with α_i greater than or equal to 3. However, $d(\nu_i, \nu_{i+1}) = 1$ and so $\iota(\nu_i, \nu_{i+1}) \leq 2$. As such, we can only have $\iota(\alpha_i, \alpha_{i+1}) \leq 2$.

II. a_{i+1} *is a wave.* Only one subcase, up to symmetry and depicted in the upper-left diagram of Figure 6, is legal. Considering only this subcase, since $d(\nu_{i+1}, \nu_{i+2}) = 1$, so the set $\{\alpha_{i+1}, \alpha'_{i+1}\} \cap \nu_{i+2}$ is non-empty. We take $j = 2$ and let $\delta_{i+2} \in \pi_Q(\{\alpha_{i+1}, \alpha'_{i+1}\} \cap \nu_{i+2})$, noting in particular that $\iota(\delta_i, \delta_{i+2}) \leq 4$ and, as such, we have $d_Q(\delta_i, \delta_{i+2}) \leq 2$.

In each of the remaining subcases, any attempt to complete the arc a_{i+1} to the curve α_{i+1} either fails or produces a curve of intersection number with α_i greater than or equal to 3.

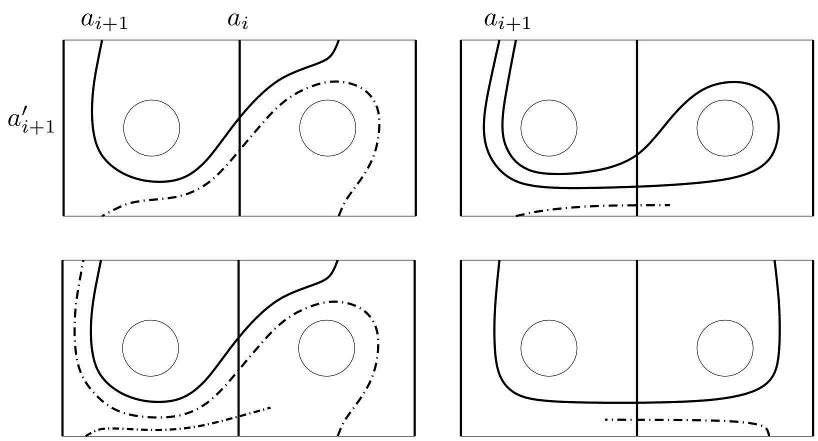

FIGURE 6. The case a_{i+1} is a wave, of which, up to symmetry, only the subcase depicted in the upper-left diagram is legal.

Acknowledgements

The authors wish to thank Brian Bowditch, Jeff Brock and Koji Fujiwara for their interest and enthusiasm, and to thank Saul Schleimer and Scott Wolpert for their suggestions on the first draft. The authors wish to thank the referee for carefully reading this paper.

The authors also wish to thank Centre Interfacultaire Bernoulli EPFL, where the first steps towards proving Theorem 1 were taken, for its hospitality and financial support. The third author also gratefully acknowledges Institute des Hautes Études Scientifiques for its hospitality and financial support. The third author is supported by a long-term Japan Society for the Promotion of Science postdoctoral fellowship, number P06034, and gratefully acknowledges the JSPS for its financial support.

References

[1] J. A. Behrstock, Y. N. Minsky, *Dimension and rank for mapping class groups*, to appear in Ann. Math. (2).
[2] J. F. Brock, *The Weil-Petersson metric and volumes of* 3*-dimensional hyperbolic convex cores*, J. Amer. Math. Soc. **16** No. 3 (2003) 495–535.
[3] J. F. Brock, B. Farb, *Curvature and rank of Teichmüller space*, Amer. J. Math. **128** (2006) 1–22.
[4] U. Hamenstädt, *Geometry of the mapping class groups III: Quasi-isometric rigidity*, arXiv:math.GT/0512429.
[5] A. E. Hatcher, W. P. Thurston, *A presentation for the mapping class group of a closed orientable surface*, Topology **19** (1980) 221–237.
[6] D. Margalit, *Automorphisms of the pants complex*, Duke Math. J. **121** No. 3 (2004) 457–479.
[7] H. A. Masur, *Extension of the Weil-Petersson metric to the boundary of Teichmuller space*, Duke Math. J. **43** no. 3 (1976) 623–635.
[8] H. A. Masur, Y. N. Minsky, *Geometry of the complex of curves II: Hierarchical structure*, Geom. Func. Anal. **10** (2000) 902–974.
[9] H. A. Masur, S. Schleimer, *The pants complex has only one end*, in "Spaces of Kleinian groups" (eds. Y. N. Minsky, M. Sakuma, C. M. Series) London Math. Soc. Lecture Note Ser. **329** (2006) 209–218.
[10] S. A. Wolpert, *Geometry of the Weil-Petersson completion of Teichmüller space*, Surveys in Differential Geometry VIII: Papers in honor of Calabi, Lawson, Siu and Uhlenbeck, editor S. T. Yau, International Press (2003).

Mathematics Institute, University of Warwick, Coventry CV4 7AL, UK
E-mail address: J.Aramayona@warwick.ac.uk
URL: http://www.maths.warwick.ac.uk/∼jaram

EPFL SB IGAT, Bâtiment BCH, CH-1015 Lausanne, Switzerland
E-mail address: hugo.parlier@a3.epfl.ch
URL: http://sma.epfl.ch/∼parlier

Dept. Mathematical and Computing Sciences, T. I. Tech., Tokyo 152-8552, Japan
E-mail address: shackleton.k.aa@m.titech.ac.jp
E-mail address: kjs2006@alumni.soton.ac.uk
URL: http://www.is.titech.ac.jp/∼kjshack5