

Copies of one-ended groups in mapping class groups

François Dahmani¹ and Koji Fujiwara²

Abstract. We establish that, given Σ a compact orientable surface and G a finitely presented one-ended group, the set of copies of G in the mapping class group $\mathcal{MCG}(\Sigma)$ consisting of only pseudo-Anosov elements except identity is finite up to conjugacy. This relies on a result of Bowditch on the same problem for images of surface groups. He asked us whether we could reduce the case of one-ended groups to his result; this is a positive answer. Our work involves analogues of Rips and Sela's canonical cylinders in curve complexes and an argument of Delzant to bound the number of images of a group in a hyperbolic group.

Mathematics Subject Classification (2000). 57M99, 20F65.

Keywords. Mapping class groups, pseudo-Anosov diffeomorphisms, curve complex.

1. Introduction

Let Σ be a compact orientable surface (possibly with boundary components). The mapping class group of Σ , denoted by $\mathcal{MCG}(\Sigma)$, is the group of isotopy classes of orientation-preserving self-homeomorphisms of Σ .

The aim of this article is to report on a control on the family of subgroups of mapping class groups that are isomorphic to a given finitely presented one-ended group.

Definition 1.1 (Purely pseudo-Anosov). A subgroup of $\mathcal{MCG}(\Sigma)$ is said purely pseudo-Anosov if all its elements, except the identity, are pseudo-Anosov mapping classes. A morphism of a group in $\mathcal{MCG}(\Sigma)$ is said purely pseudo-Anosov if it is injective and has purely pseudo-Anosov image.

Up to now, the only known purely pseudo-Anosov subgroups of mapping class groups are free.

Recently, Brian Bowditch [4] has established the finiteness of the set of such images of surface groups, up to conjugacy. He uses deep results, in particular in

¹The first author acknowledges partial support from the ANR grant ANR-06-JCJC-0099.

²The second author is partially supported by Grant-in-Aid for Scientific Research (No. 19340013).

3-manifold geometry, and doing so, he proves a general powerful result [4], Proposition 8.1. He asked us, however, whether it is possible to adapt the situation of an arbitrary finitely presented one-ended group to the setting of his proposition. We provide here an affirmative answer.

Theorem 1.2. *Given Σ a compact orientable surface and G a finitely presented one-ended group, there is a number N such that any purely pseudo-Anosov subgroup of $\mathcal{MCG}(\Sigma)$ isomorphic to G admits a set of generators (γ_i) for which there are a vertex v in the curve complex $\mathcal{CC}(\Sigma)$ with $d(\gamma_i v, v) \leq N$, and a presentation on this set of generators with at most N relators, each of them of length at most N as words.*

This allows to apply the following important result of Bowditch.

Proposition 1.3 ([4], Proposition 8.1). *Suppose that G is a one-ended finitely presented group and that $\phi: G \rightarrow \mathcal{MCG}(\Sigma)$ is a purely pseudo-Anosov homomorphism, giving an induced action on $\mathcal{CC}(\Sigma)$. Suppose that A is a generating set of G and that there is some $v \in \mathcal{CC}(\Sigma)$ and $N > 0$ such that $d(\phi(a)v, v) \leq N$ for all $a \in A$. Then there is some $\theta \in \mathcal{MCG}(\Sigma)$ such that the word length of each $\theta\phi(a)\theta^{-1}$, $a \in A$ (in terms of a generating set of $\mathcal{MCG}(\Sigma)$), is bounded above in terms of N and the sum of lengths of relators in a presentation of G on the generating set A .*

The following corollary is immediate, since for each given G and A the set of elements $\theta\phi(a)\theta^{-1}$, $a \in A$, is finite for all ϕ if we choose θ using the proposition for each ϕ .

Corollary 1.4. *Given Σ a compact orientable surface and G a finitely presented one-ended group, the set of purely pseudo-Anosov subgroups of $\mathcal{MCG}(\Sigma)$ isomorphic to G is finite up to conjugacy in $\mathcal{MCG}(\Sigma)$.*

The group $\mathcal{MCG}(\Sigma)$ has a natural action by isometries on Harvey's curve complex $\mathcal{CC}(\Sigma)$, which turns out to be a hyperbolic space [8], [2]. This complex is far from being locally finite, and the action is not proper.

The elements of $\mathcal{MCG}(\Sigma)$ that are hyperbolic isometries of $\mathcal{CC}(\Sigma)$ are precisely the pseudo-Anosov elements of $\mathcal{MCG}(\Sigma)$.

Our method towards Theorem 1.2 is inspired by the case of relative hyperbolicity studied in [6] (and indeed by the hyperbolic case, [7]): construct Rips and Sela's canonical cylinders, as in [9], for a given morphism $G \rightarrow \mathcal{MCG}(\Sigma)$, and use them to pull back a lamination on a van Kampen complex $P(G)$ of G (or more precisely, first on its universal cover) that allows to find small generators of the image.

To perform the construction, we make use of deep results of Bowditch about tight geodesics in $\mathcal{CC}(\Sigma)$.

In this article we introduce the relevant definitions for the general argument but sometimes refer to existing proofs when they can be applied without modification.

We tried to make clear where the technology specific to mapping class groups is used, or where the existing argument would not, as written, give sufficient precision. In particular, our main task about Theorem 2.9, which is based on very subtle ideas of Rips and Sela, is to explain how to get a setting where the original proof can be applied (which is not obvious without Bowditch's results). However, for the reader's convenience, we also reproduce this proof in Section 2.1.

Let us mention that in the case that G is also a surface group other proofs of Theorem 1.2 have been given, notably by J. Barnard [1].

We thank B. Bowditch for stimulating discussions and for asking the question on the bound on the complexity of presentations. We learned, while finishing this article, that he very simultaneously obtained a similar result, by different methods, using actions on \mathbb{R} -trees [5]. The first author thanks T. Delzant for interesting and stimulating related discussions. Both authors want to thank the referee for constructive remarks. The second author gratefully acknowledges the Institut de Mathématiques de Toulouse and support from CNRS. This work was partially done while he was visiting the institute.

2. Sliced canonical cylinders

In the following, $\mathcal{CG}(\Sigma)$ is the curve graph of a surface Σ (the one-skeleton of Harvey's curve complex $\mathcal{CC}(\Sigma)$), $\delta \in \mathbb{N}$ is an hyperbolicity constant and p is a base point in $\mathcal{CG}(\Sigma)$.

The graph $\mathcal{CG}(\Sigma)$ is not locally finite, and in general two points are joined by infinitely many different geodesics. However, there is a class of geodesics that are called *tight geodesics*, with good properties. We will not need the definition, which involves properties of the curves and subsurfaces in the surface Σ , so we just refer to [3] for it. We will need the fact that there exist such geodesics, that they satisfy the statement of Theorem 2.5, and that a subpath of a tight geodesic is a tight geodesic.

Definition 2.1 (λ -quasi-geodesic). A λ -quasi-geodesic in a graph X is a λ -bi-lipschitz embedding of a segment of \mathbb{R} into X . We assume here that paths start and end on vertices. The length of a path is the number of edges in its image.

A path is a ν -local- λ -quasi-geodesic if any of its subpaths of length at most ν is a λ -quasi-geodesic.

A path is a μ -local tight geodesic if any of its subpaths of length μ is a tight geodesic.

We choose some constants: $\lambda = 1000\delta$ and ε such that any λ -quasi geodesic in $\mathcal{CG}(\Sigma)$ stays ε -close to any geodesic joining its end points. Let $\mu = (100\varepsilon + \lambda^2) \times 40\lambda$ and $\nu = 40\lambda(\varepsilon + 100\lambda\delta)$.

The next definition is to be compared with a similar one in [9], for geodesics that are not necessarily tight.

Definition 2.2 (Coarse piecewise tight geodesics or *cptg*). Let $l \geq \mu$ be an integer. An l -coarse piecewise tight geodesic, or l -*cptg*, in $\mathcal{CG}(\Sigma)$ is a ν -local $\frac{\lambda}{2}$ -quasigeodesic $f: [a, b] \rightarrow \mathcal{CG}(\Sigma)$ together with a subdivision $a = c_1 \leq d_1 \leq c_2 \leq \dots \leq d_n = b$ of the segment $[a, b]$ such that $f([c_i, d_i])$ is a μ -local tight geodesic of length at least l when $2 \leq i \leq (n - 1)$ and $\text{length}(f[d_i, c_{i+1}]) \leq \varepsilon$ for all $i \leq n$.

Moreover we require $f([a, b])$ be in the 2ε -neighborhood of a tight geodesic segment $[f(a), f(b)]$.

Remark. Any l -*cptg* is a λ -quasi-geodesic (this does not use tightness but only hyperbolic geometry; see for instance [6], Appendix). The subpaths corresponding to a subdivision of an l -*cptg* as in the definition are called sub-local geodesics and bridges, respectively.

Definition 2.3 (l -Cylinders, [9]). Let $l \in \mathbb{N}$. The l -cylinder of two points x and y in $\mathcal{CG}(\Sigma)$, denoted by $\text{Cyl}_l(x, y)$, is the set of the vertices v lying on an l -*cptg* from x to y , with the additional requirement that v is on a local tight geodesic $f|_{[c,d]}$ of the subdivision, with distances $|f(c) - v| \geq l$ if $f(c) \neq x$ and $|f(d) - v| \geq l$ if $f(d) \neq y$.

Any tight geodesic is an l -*cptg* with a trivial subdivision and is thus contained in the l -cylinder between its end points for all l . Here is an obvious consequence of the definitions.

Lemma 2.4 (Equivariance). $\gamma \text{Cyl}_l(x, y) = \text{Cyl}_l(\gamma x, \gamma y)$ for all $\gamma \in \mathcal{MC}\mathcal{G}(\Sigma)$ and for all x, y . Moreover, $\text{Cyl}_l(x, y) = \text{Cyl}_l(y, x)$.

Recall a crucial result of Bowditch (that will be used in the next two lemmas).

Theorem 2.5 ([3], Theorem 1.1 and Theorem 1.2). Given L , there is a constant $K_0(L)$ such that, for all a, b , vertices of the curve complex and all c on a tight geodesic between a and b , the set of vertices on tight geodesics between a and b and at distance at most $L + 2\varepsilon$ from c has at most $K_0(L)$ elements.

There are constants k_1 and K_1 depending only on Σ such that if a, b are vertices in $\mathcal{CC}(\Sigma)$, $r \in \mathbb{N}$, and c on a tight geodesic joining a to b , with $d(a, c) \geq r + k_1$ and $d(b, c) \geq r + k_1$, then the set of vertices on tight geodesics between two points respectively r -close of a and b , and at distance at most 2ε from c , has at most K_1 elements.

Lemma 2.6. Let $l \geq k_1 + \mu/2$. Any l -cylinder of the curve complex $\mathcal{CC}(\Sigma)$ is finite.

Proof. Let x, y be two vertices in the curve complex. Let C be the set of vertices on tight geodesics between x, y that are at distance at least $l - \mu/4$ from both x and y . By the first point of Theorem 2.5, this set is finite. For each $v \in C$ let B_v be the set given by the second point of Theorem 2.5 for $r = \mu/4$ (which can be applied since l

is large enough). Let B be the union of all the B_v , $v \in C$; this set is finite. Let also B' be the set of all vertices on tight geodesics from x to y or from x to a vertex of B , or from y to a vertex of B . Since B is finite and by the first point of Theorem 2.5, this set B' is finite. We want to show that the cylinder of x, y is a subset of $B \cup B'$.

Let w lie in an l -cylinder of x, y . It is on a local tight geodesic $f|_{[c,d]}$ with $f(c)$ and $f(d)$ at distance at most 2ε from a tight geodesic between x and y .

Thus, w is on a subsegment σ of length μ that is a tight geodesic and whose end points are at distance at most 2ε from a tight geodesic $[x, y]$. There are two cases following from the inequality condition of Definition 2.3: either we can assume that one of the ends of σ is x (or y) and $d(w, x) \leq \mu/2$ (or similarly with y), or we can assume that w is in the middle of σ .

In the second case, since $\mu > 100\varepsilon$, one can find a smaller sub-(tight geodesic) of length $\mu/2$ with w in its middle and whose end points are, by hyperbolicity, at distance at most $2\delta (< \mu/4)$ from a tight geodesic $[x, y]$. Then by definition of the sets B_v , we have that w in some set B_v for some $v \in C$.

In the first case, assume that one end of σ is x . Then applying the above argument to the center $c(\sigma)$ of σ , we find that $c(\sigma) \in B$, and therefore $w \in B'$. \square

Definition 2.7 (Channels, cf. [9], 4.1). Let $L > 0$ and let $a, b \in \mathcal{CG}(\Sigma)$ with $d(a, b) \leq 3L$. A tight geodesic g_1 of length L which is contained in a tight geodesic g_2 of length $3L$ that starts (respectively ends) at distance at most 2ε from a (respectively b), such that end points of g_2 are at distance L from g_1 , is called an L -channel of (a, b) .

Lemma 2.8. *For every $L \geq k_1 + 4\varepsilon$, there is a bound $\kappa(L)$ on the number of L -channels of (a, b) for arbitrary a, b with $d(a, b) \leq 3L$.*

Proof. First we can assume that $d(a, b) \geq 3L - 4\varepsilon$; otherwise there are no channels at all.

We will show a finite set (of cardinality uniformly bounded above) containing all vertices of L -channels of a, b . Let C be the set of vertices on tight geodesics from a to b that are at distance at least $L - 2\varepsilon$ from both. Because $d(a, b) \leq 3L$, by the first point of Theorem 2.5, C has at most $K_0(L)$ elements. For each $c \in C$, consider the set B_c given by the second point of Bowditch's theorem, for $r = 2\varepsilon$, applied to c (which is justified since $d(a, c) \geq L - 2\varepsilon \geq k_1 + 2\varepsilon$ by choice of L). Let B be the union of all the B_c , $c \in C$. It has at most $K_0 \times K_1$ elements.

Now consider w a vertex on an L -channel g_1 of a, b , which by definition is a sub-(tight geodesic) of a tight geodesic g_2 starting and ending at a', b' with $d(a, a') \leq 2\varepsilon$ and $d(b, b') \leq 2\varepsilon$. Let $L + x = d(a', w)$ ($x \geq 0$), and assume (by symmetry this is without loss of generality) that $x \leq L/2$. We have $L - 2\varepsilon + x \leq d(a, w) \leq L + 2\varepsilon + x$. By hyperbolicity, w is 2δ -close to a point w' in a segment $[a, b]$ (which we choose tight). Since $2\delta \leq \varepsilon$, we can find another point w'' on $[a, b]$ at distance at most ε from w (hence $w \in B_{w''}$) and at least $L - 2\varepsilon + x$ from a (hence $w'' \in C$). This shows that $w \in B$. \square

For an integer n , we set $\psi(n) = 24(n + 1)\kappa(\mu)(2\varepsilon + 1)\varepsilon$. We denote by $B_r(x)$ the ball of $\mathcal{C}\mathcal{G}(\Sigma)$ of center x and radius r .

Theorem 2.9. *Let F be a finite family of elements of $\mathcal{M}\mathcal{C}\mathcal{G}$; we set $n = (2\sharp(F))^3$ where $\sharp(F)$ is the cardinality of F . Let p be a base point in $\mathcal{C}\mathcal{G}(\Sigma)$.*

There exists a number l such that the l -cylinders satisfy: for all α, β, γ in $F \cup F^{-1}$ with $\alpha\beta\gamma = 1$, in the triangle $(x, y, z) = (p, \alpha p, \gamma^{-1} p)$ in $\mathcal{C}\mathcal{G}(\Sigma)$ one has

$$\text{Cyl}_l(x, y) \cap B_{R_{x,y,z}}(x) = \text{Cyl}_l(x, z) \cap B_{R_{x,y,z}}(x)$$

(and analogues permuting x, y and z) where $R_{x,y,z} = (y \cdot z)_x - 5 \times (13\mu + \psi(n))$ is the Gromov product in the triangle, minus a constant.

Note that in the statement, l depends on F , but $(y \cdot z)_x - R_{x,y,z}$ depends only on $\sharp(F)$.

One may think that $\text{Cyl}_l(x, y)$ is a narrow set near a geodesic from x to y . The theorem says that $\text{Cyl}_l(x, y), \text{Cyl}_l(y, z)$ and $\text{Cyl}_l(z, x)$ coincide except in a set of bounded size near the center of the triangle (x, y, z) . Instead of $\text{Cyl}_l(x, y)$, if we take the union of all (tight) geodesics between x, y or the union of all quasi-geodesics between x, y with uniform quasi-geodesic constants, we do not have this equation in general. This is already the case for a Cayley graph of a word-hyperbolic group, and Rips–Sela [9] introduced several notions in this context which we imitated here.

2.1. Proof of Theorem 2.9. We produced a setting where cylinders and channels are finite. We can therefore reproduce the original proof of Theorem 2.9 by Rips and Sela for hyperbolic groups [9]. For the reader’s convenience, we give the detail. We follow the exposition in [6], Theorem 2.9 (which was for relatively hyperbolic groups).

Let us start by stating two lemmas for rerouting a cptg.

Lemma 2.10. *Let $l \geq \mu$ and $f : [a, b] \rightarrow \mathcal{C}\mathcal{G}$ be an l -cptg, whose subdivision includes $f|_{[c,d]}$, a local tight geodesic. Let $s = f(t) \in f([c, d])$ such that the path $f([c, t])$ from $f(c)$ to s has length $\geq l + 2\varepsilon$.*

Let now g be a tight geodesic segment joining $f(a)$ to $f(b)$ and let s'' be a point on g closest to s . Let $s' = f(t') \in f([c, d])$ be a point closest to s'' on $f([c, d])$.

Let $[s', s'']$ be a geodesic segment and $[s'', f(b)]$ be a subsegment of g .

*Then the concatenation $f([a, t']) * [s', s''] * [s'', f(b)]$ is an l -cptg from $f(a)$ to $f(b)$, and we say that f can be rerouted into this new path.*

Lemma 2.11. *Let f be an l -cptg whose last subdivision segment is a local tight geodesic g of length at least $l + 2\mu$. Let $z \in \mathcal{C}\mathcal{G}$ such that a tight geodesic segment $[f(a), z]$ passes at distance at most δ from $f(b)$.*

Then there exists an l -cptg from $f(a)$ to z coinciding with f until the first point of g .

We do not repeat the proofs of these two lemmas here. They are rather standard, we refer, for instance, to Lemmas 2.2 and 2.4 in [6]; see Figure 1 for an illustration. The main observation is that the proposed paths are indeed local quasi-geodesics (the other properties being immediate).

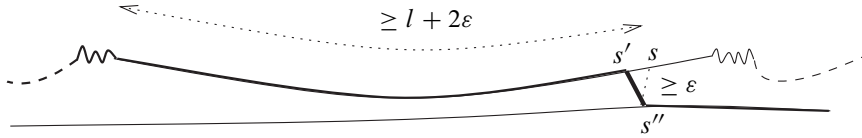


Figure 1. A typical example of rerouting a cptg (see Lemmas 2.10 and 2.11). Starting with a cptg one of whose local geodesics is longer than $l + 2\varepsilon$, one deduces another cptg (bold in the picture) coinciding with the first until a certain point on this long local geodesic (s' on the picture), but ending to a possibly different point.

Lemma 2.12. *For all integers i in $\{1, \dots, \frac{\psi(n)}{2\varepsilon}\}$, put $l_i = 10\mu + 2i\varepsilon \leq \psi(n) + 10\mu$. Let x, y, z be three points in $\mathcal{C}\mathcal{G}$. There are at most $2\kappa(\mu) \times (2\varepsilon + 1)$ different values of l_i for $i \in \{1, \dots, \frac{\psi(n)}{2\varepsilon}\}$ such that*

$$\text{Cyl}_{l_i}(x, y) \cap B_{R_{x,y,z}}(x) \not\subset \text{Cyl}_{l_i}(x, z) \cap B_{R_{x,y,z}}(x).$$

Proof. We argue by contradiction, assuming that $2\kappa(\mu) \times (2\varepsilon + 1) + 1$ different l_i do not satisfy $\text{Cyl}_{l_i}(x, y) \cap B_{R_{x,y,z}}(x) \subset \text{Cyl}_{l_i}(x, z) \cap B_{R_{x,y,z}}(x)$.

For each of them, there is $v_i \in B_{R_{x,y,z}}(x)$ in $\text{Cyl}_{l_i}(x, y)$ but not in $\text{Cyl}_{l_i}(x, z)$: there is an l_i -cptg, β_i , from x to y containing v_i as indicated in Definition 2.3, and none from x to z .

By definition of cptg, given a tight geodesic $[x, y]$, β_i is contained in its 2ε -neighborhood. Thus every subsegment of length μ of a sub-local geodesic of β_i , at distance at least μ from the end points of the sub-local geodesic, is in fact a μ -channel of some subsegment of $[x, y]$.

Let $[x', y']$ and $[x'', y'']$ be two subsegments of $[x, y]$ of length 3μ such that $d(x', x) = R_{x,y,z} + (\psi(n) + 9\mu)$ and $d(x'', x) = R_{x,y,z} + (\psi(n) + 13\mu)$ (the end of the first of these segments is at distance μ from the beginning of the second). Assume that β_i does not contain a μ -channel of $[x', y']$. This means, as we just noticed, that it must have a bridge $\beta_i([d_j, c_{j+1}])$ at distance $3\mu + 2\varepsilon$ from x' . Since $l_i \geq 10\mu$ (and $\varepsilon \leq \mu/10$), the sub-local geodesic after this bridge must contain a μ -channel of $[x'', y'']$. Therefore, each β_i contains a μ -channel of either $[x', y']$ or $[x'', y'']$.

There are at most $2\kappa(\mu)$ different μ -channels of either $[x', y']$ or $[x'', y'']$. Hence there is a channel, denoted by Chan, through which some local geodesic subdivision $\beta_i|_{[c(i), d(i)]}$ passes for at least $2\varepsilon + 2$ different indices i . Let $i_1 < \dots < i_{2\varepsilon+2}$ be such indices.

For each j , let $t_{i_j} \in [c(i_j), d(i_j)]$ be the instant where $\beta_{i_j}(t_{i_j})$ exists the channel Chan, and r_{i_j} the length of the path $\beta_{i_j}([t_{i_j}, d(i_j)])$. There are three claims about the possible values of r_{i_j} .

Claim 1: For any $j \in \{1, \dots, 2\varepsilon + 2\}$, one has $r_{i_j} < l_{i_j} + 2\varepsilon$.

Claim 2: If $i_j < i_k$ then $r_{i_k} < r_{i_j}$.

Claim 3: $r_{i_1} - r_{i_{2\varepsilon+2}} < 2\varepsilon$.

From the second claim, we deduce that all the r_{i_j} are different, and from the third claim we deduce that they are integers in an interval of length $2\varepsilon + 1$. Since there are $2\varepsilon + 2$ values this is a contradiction.

We now have to prove these claims.

For the first one, assume the contrary, and let $t_{i_j}^+ > t_{i_j}$ be a real number such that the length of $\beta_{i_j}([t_{i_j}, t_{i_j}^+])$ is l_{i_j} . Our assumption allows to use Lemma 2.10: one can change β_{i_j} into another l_{i_j} -cptg coinciding with $[x, y]$ on a subsegment containing $[y', y]$, with $d(y', \beta_{i_j}(t_{i_j}^+)) \leq 3\varepsilon$. By the triangle inequality we have $d(x, \beta_{i_j}(t_{i_j})) \leq R_{x,y,z} + (\psi(n) + 13\mu) + 2\varepsilon + 2\mu$. Therefore it follows that $d(x, \beta_{i_j}(t_{i_j}^+)) \leq R_{x,y,z} + (\psi(n) + 13\mu) + 2\varepsilon + 2\mu + (\psi(n) + 10\mu) + 2\varepsilon$, which is $\leq R_{x,y,z} + 2(\psi(n) + 13\mu) \leq (z \cdot y)_x - 2(\psi(n) + 13\mu)$, meaning that it is at least $l_i + 2\mu$ before reaching a point δ -close to the center of the triangle (x, y, z) . This allows to use Lemma 2.11: the new cptg can be rerouted into another one, for the same l_{i_j} , coinciding with the beginning β_{i_j} until after v_{i_j} , and ending at z . In particular, v_{i_j} is in $\text{Cyl}_{l_{i_j}}(x, z)$, contradicting our assumption.

Now we use the first claim to prove the second. If the latter were not true, one could change β_{i_j} into β_{i_k} just after its passage in Chan (it is enough to notice that this new path remains an l_{i_j} -cptg since $i_j < i_k$). On β_{i_k} , consider the sub-local tight geodesic of the subdivision following that of Chan. Because $i_k > i_j$, it is longer than $l_{i_j} + 2\varepsilon$; let $\beta_{i_k}(t_{i_k}^+)$ be the point on it after travelling this distance (which is $\leq (\psi(n) + 10\mu) + 2\varepsilon$ in any case). As before, $d(x, \beta_{i_j}(t_{i_j})) \leq R_{x,y,z} + (\psi(n) + 13\mu) + 2\varepsilon + 2\mu$. By the first claim, $r_{i_j} \leq \psi(n) + 10\mu + 2\varepsilon$, then the next bridge is at most ε long, and we need to travel at most $(\psi(n) + 10\mu) + 2\varepsilon$ further to find $\beta_{i_j}(t_{i_j}^+)$. Thus, $d(x, \beta_{i_j}(t_{i_j}^+)) \leq R_{x,y,z} + (\psi(n) + 13\mu) + 2\varepsilon + 2\mu + 2(\psi(n) + 10\mu) + 5\varepsilon$, which is $\leq R_{x,y,z} + 3(\psi(n) + 13\mu) \leq (z \cdot y)_x - 2(\psi(n) + 13\mu)$. We then use, as in claim 1, Lemma 2.10 and Lemma 2.11 to obtain the same contradiction.

The third claim is again proved by contradiction: if it were false, we could change $\beta_{i_{2\varepsilon+2k}}$ just after Chan by substituting the remaining part of the sub-local tight geodesic of β_{i_1} containing Chan. Then one can reroute this cptg on $[x, y]$ at distance 2ε before the end of this sub-local geodesic, and finally reroute it again into a cptg ending at z , again a contradiction. \square

Now we can prove Theorem 2.9.

Proof of Theorem 2.9. We need to find a good parameter l . For this we have at

least $\frac{\psi(n)}{2^\varepsilon} = 12(n + 1)\kappa(\mu) \times (2\varepsilon + 1)$ candidates: the parameters l_i defined in Lemma 2.12. There are at most n different triangles satisfying the condition of Theorem 2.9. Therefore we have a system of at most $6n$ inclusions of the form $\text{Cyl}_{l_i}(x, y) \cap B_{R_{x,y,z}}(x) \subset \text{Cyl}_{l_i}(x, z) \cap B_{R_{x,y,z}}(x)$ to satisfy simultaneously. For each inclusion, by Lemma 2.12, only $2\kappa(\mu)(2\varepsilon + 1)$ parameters l_i fail to satisfy it. Hence, by the pigeonhole principle, one parameter satisfies all the $6n$ inclusions. \square

2.2. Slicing. Let us assume that l satisfies the conclusion of Theorem 2.9. From now on, all cylinders will be l -cylinders, and we write $\text{Cyl}(a, b)$ for $\text{Cyl}_l(a, b)$.

Let $\text{Cyl}(a, b)$ be a cylinder and $x \in \text{Cyl}(a, b)$. Following [9], we define the set $N_{\mathbf{R}}^{(a,b)}(x)$ as follows: it is the set of all the vertices $v \in \text{Cyl}(a, b)$ such that $|a - x| < |a - v|$ and $|x - v| > 100\delta$. Here “ \mathbf{R} ” stands for “right”, and $N_{\mathbf{L}}^{(a,b)}(x)$ is similarly defined changing the condition $|a - x| < |a - v|$ into $|a - x| > |a - v|$. As cylinders are finite, these sets are also finite.

Let $x, y \in \text{Cyl}(a, b)$ be in a cylinder. We set

$$\begin{aligned} \text{Diff}_{a,b}(x, y) = & \#(N_{\mathbf{L}}^{(a,b)}(x) \setminus N_{\mathbf{L}}^{(a,b)}(y)) - \#(N_{\mathbf{L}}^{(a,b)}(y) \setminus N_{\mathbf{L}}^{(a,b)}(x)) \\ & + \#(N_{\mathbf{R}}^{(a,b)}(y) \setminus N_{\mathbf{R}}^{(a,b)}(x)) - \#(N_{\mathbf{R}}^{(a,b)}(x) \setminus N_{\mathbf{R}}^{(a,b)}(y)), \end{aligned}$$

where $\#(X)$ is the cardinality of the set X . This definition makes sense since because of Lemma 2.6 all the sets involved are finite.

Lemma 2.13. *Diff_{a,b} satisfies a cocycle relation: for arbitrary $x, y, z \in \text{Cyl}(a, b)$, one has $\text{Diff}_{a,b}(x, z) = \text{Diff}_{a,b}(x, y) + \text{Diff}_{a,b}(y, z)$. In particular, the relation $\text{Diff}_{a,b}(x, y) = 0$ is an equivalence relation on $\text{Cyl}(a, b)$.*

Let us say that an equivalence class for this relation ($\text{Diff}_{a,b}(x, y) = 0$) is a slice of $\text{Cyl}(a, b)$. Then the value of $\text{Diff}_{a,b}(x, y)$ depends only on the slices of x and y . Moreover, the relation on the set of slices defined by $S < S'$ if $\text{Diff}_{a,b}(x, y) < 0$ for all $x \in S, y \in S'$ is a total order on the set of slices.

Proof. All the assertions are immediate consequences of the first one, which follows from a short computation (we reproduce that of [9], Lemma 3.4). Notice that $\#(N_{\mathbf{L}}(x) \setminus N_{\mathbf{L}}(y)) - \#(N_{\mathbf{L}}(y) \setminus N_{\mathbf{L}}(x)) + \#(N_{\mathbf{L}}(y) \setminus N_{\mathbf{L}}(z)) - \#(N_{\mathbf{L}}(z) \setminus N_{\mathbf{L}}(y))$ is equal to $\#(N_{\mathbf{L}}(x) \setminus N_{\mathbf{L}}(z)) - \#(N_{\mathbf{L}}(z) \setminus N_{\mathbf{L}}(x))$, where we write N for $N^{a,b}$, and similarly for $N_{\mathbf{R}}$. \square

Lemma 2.14 (Properties of slices). (i) *If $v \in \text{Cyl}(a, b)$, then v is at distance at most 2δ from any tight geodesic segment $[a, b]$.*

(ii) *Let S be a slice of $\text{Cyl}(a, b)$ and let v, v' be in S . Then $d(v, v') \leq 200\delta$.*

(iii) *Let v and v' be in two consecutive slices of $\text{Cyl}(a, b)$. Then $|v - v'| \leq 1000\delta$.*

(iv) *If $\text{Cyl}(a, b) \cap B_R(a) = \text{Cyl}(a, c) \cap B_R(a)$, where $B_R(a)$ is the ball of radius R centered at a , then any slice of $\text{Cyl}(a, b)$ included in $B_{R-200\delta}(a)$ is a slice of $\text{Cyl}(a, c)$.*

Proof. The following is a repetition of the proofs of Lemmas 2.19–2.21 from [6].

For (i) it suffices to see that a point in a cylinder is in a geodesic starting and ending at distance 2ε from $[a, b]$, and sufficiently far from its endpoints.

For (ii) assume that $d(a, v) \leq d(a, v')$ and $d(v, v') \geq 200\delta$. Then the result follows from the relations $N_L^{(a,b)}(v) \subset N_L^{(a,b)}(v')$ (strict inclusion) and $N_R^{(a,b)}(v') \subset N_R^{(a,b)}(v)$. We only show the first one (the second one is similar). The equality is impossible since v' is in one set and not in the other. Let $w \in [a, b]$ at distance 2δ from v , and similarly $w' \in [a, b]$ close to v' . Clearly $d(a, w) \leq d(a, w') - 196\delta$. If $z \in N_L^{(a,b)}(v)$, there is w_z on $[a, b]$ at distance 2δ from it. By definition of $N_L^{(a,b)}(v)$ it follows that $d(w, w_z) \geq 96\delta$, and $w_z \in [a, w]$. By the triangle inequality, one finds that $z \in N_L^{(a,b)}(v')$, which is what we wanted.

For (iii) assume that $d(a, v) < d(a, v')$ and $d(v, v') > 1000\delta$. We can find $w \in [a, b]$ at distance at least 400δ from v and from v' and such that $d(a, v) + 200\delta < d(a, w) < d(a, v') - 200\delta$. It is easy to check that w is in a slice between that of v and that of v' , contradicting that they are in consecutive slices.

For (iv) let $S, S' \subset B_{R-200\delta}(a)$ be slices of $\text{Cyl}(a, b)$ and $\text{Cyl}(a, c)$, respectively. We claim that if they intersect, they are equal. Let $v \in S \cap S'$ and $v' \in S$. It is enough to check that $\text{Diff}_{a,b}(v, v') = \text{Diff}_{a,c}(v, v')$, because this would imply that $v' \in S'$ and $S \subset S'$, and by symmetry, equality holds. By assumption we have $\text{Cyl}(a, b) \cap B_R(a) = \text{Cyl}(a, c) \cap B_R(a)$. Therefore $N_L^{(a,c)}(v) = N_L^{(a,b)}(v)$, and similarly for v' . If $x \in N_R^{(a,c)}(v') \setminus N_R^{(a,c)}(v)$, it is 100δ -close to v , and it is then in $\text{Cyl}(a, b)$ and in $N_R^{(a,b)}(v') \setminus N_R^{(a,b)}(v)$. By symmetry we also have the reverse inclusion, and $N_R^{(a,c)}(v') \setminus N_R^{(a,c)}(v) = N_R^{(a,b)}(v') \setminus N_R^{(a,b)}(v)$, which ensures that $\text{Diff}_{a,b}(v, v') = \text{Diff}_{a,c}(v, v')$. \square

As a consequence of Theorem 2.9 and Lemma 2.14 (iv), one obtains:

Proposition 2.15. *With the notations of Theorem 2.9, let l be the constant given by it. Let $(x, y, z) = (p, \alpha p, \gamma^{-1} p)$ be a triangle in $\mathcal{CE}(\Sigma)$ such that α, β, γ are in $F \cup F^{-1}$ and $\alpha\beta\gamma = 1$.*

The ordered slice decomposition of the cylinders is

$$\begin{aligned} \text{Cyl}_l(x, y) &= (S_1, S_2, \dots, S_k, \mathcal{H}_z, T_m, T_{m-1}, \dots, T_1), \\ \text{Cyl}_l(x, z) &= (S_1, S_2, \dots, S_k, \mathcal{H}_y, V_p, V_{p-1}, \dots, V_1), \\ \text{Cyl}_l(y, z) &= (T_1, T_2, \dots, T_m, \mathcal{H}_x, V_p, V_{p-1}, \dots, V_1), \end{aligned}$$

such that $S_1, \dots, S_k, T_1, \dots, T_m$ and V_1, \dots, V_p are slices and each $\mathcal{H}_v, v = x, y, z$, is a set of at most $10\psi(n)$ consecutive slices. The sets \mathcal{H}_v are called the holes of the slice decomposition.

3. Purely pseudo-Anosov images of groups

To obtain Theorem 1.2, we can follow the approach in [7] without major change.

Let us consider $\varphi: G \rightarrow \mathcal{MC}\mathcal{G}(\Sigma)$ and a van Kampen 2-complex $P(G)$ of G (so $G \simeq \pi_1(P(G))$) once a base point is chosen). $P(G)$ is a simplicial complex with one vertex (and base point), a certain number of edges e_1, \dots, e_r that we identify with elements of G , and $T(G)$ triangles (which are the relators of a triangular presentation).¹ Then we pass to the universal cover $\widetilde{P(G)}$, where we choose a base point and representatives \tilde{e}_i of the e_i starting at this point. The vertices of $\widetilde{P(G)}$ are thus identified with the group $\pi_1(P(G))$, and the map φ induces a map from the vertices of $\widetilde{P(G)}$ to $\mathcal{MC}\mathcal{G}(\Sigma)$, and therefore, by considering the orbit of the base point $p \in \mathcal{C}\mathcal{G}(\Sigma)$, to the curve graph $\mathcal{C}\mathcal{G}(\Sigma)$. Denote this map by $\tilde{\varphi}: \widetilde{P(G)}^{(0)} \rightarrow \mathcal{C}\mathcal{G}(\Sigma)$. We apply Proposition 2.15 to the family $\varphi(e_1), \dots, \varphi(e_r)$, thus providing canonical cylinders for each pair $(p, \tilde{\varphi}(e_i)p)$ in $\mathcal{C}\mathcal{G}(\Sigma)$ and their translates (we will omit the constant l in notations, since it is now fixed until the end). We extend the map $\tilde{\varphi}$ to the 1-skeleton of $\widetilde{P(G)}$ by mapping, for all i , the edge \tilde{e}_i onto a path in $\text{Cyl}(p, \varphi(e_i)p)$ that successively goes through all consecutive slices, and then extend the map $\tilde{\varphi}$ on the translates of \tilde{e}_i equivariantly. This gives an equivariant map $\tilde{\varphi}: \widetilde{P(G)}^{(1)} \rightarrow \mathcal{C}\mathcal{G}(\Sigma)$.

If $\text{Cyl}(p, \varphi(e_i)p)$ has n slices, we choose n distinct points $m_k, k = 1, \dots, n$ (we call them “marked points”), on the edge \tilde{e}_i such that $\tilde{\varphi}(m_k)$ is in the k -th slice of $\text{Cyl}(p, \varphi(e_i)p)$ for all k . We complete by translation, so that every edge $\tilde{e} = (v_1, v_2)$ of $\widetilde{S_g}$ has a certain number of points marked on it that are mapped into the consecutive slices of the cylinder $\text{Cyl}(\tilde{\varphi}(v_1), \tilde{\varphi}(v_2))$. We will use the coincidence of slices to construct tracks in $P(G)$.

Let us consider a representative of an orbit of triangular cells in $\widetilde{P(G)}$. We link each pair of marked points on the edges by a “blue” segment, when the slices in which they are mapped are equal. After that, in each triangle where there are unlinked marked points, we add a singular red point in the triangle and link it with every remaining marked point by red edges that do not cross any blue one (it is clear that there is a way of choosing the red point so that this is possible). By Proposition 2.15, each singular red point is linked by a red edge to at most $30\psi(T(G))$ marked points.

The union \mathcal{BR} of these segments defines a family of disjoint connected graphs in $\widetilde{P(G)}$, some graphs with blue and red edges, and some graphs with only blue edges.

We now extend the map $\tilde{\varphi}$ on each of these graphs. It suffices to choose the image of each blue or red segment joining two marked points in a triangle, and then complete by translations. We thus choose any path in $\mathcal{C}\mathcal{G}(\Sigma)$ joining the images of its end points (and similarly for red ones). By construction this map $\tilde{\varphi}: \widetilde{P(G)}^{(1)} \cup \mathcal{BR} \rightarrow \mathcal{C}\mathcal{G}(\Sigma)$ is still equivariant, meaning that $\tilde{\varphi}(gv) = \varphi(g)\tilde{\varphi}(v)$ for all $v \in \widetilde{P(G)}$ and all $g \in G$.

¹In principle we could use relations of length 2 or 3, but since pseudo-Anosov elements of $\mathcal{MC}\mathcal{G}$ are of infinite order, we can assume G to be torsion-free, in particular without element of order 2, and it is not hard then to eliminate relations of length 2.

Hence we have:

Lemma 3.1. *The image of a connected completely blue graph of \mathcal{BR} is contained in a single slice, and thus is finite.*

Since the construction was done G -equivariantly, \mathcal{BR} descends to the quotient $P(G)$ as $\overline{\mathcal{BR}}$, which denotes the union of disjoint connected graphs, some of them completely blue, some of them containing red edges. Since there are $T(G)$ triangles in $P(G)$, there are at most $T(G) \times 30\psi(T(G))$ red edges in $\overline{\mathcal{BR}}$.

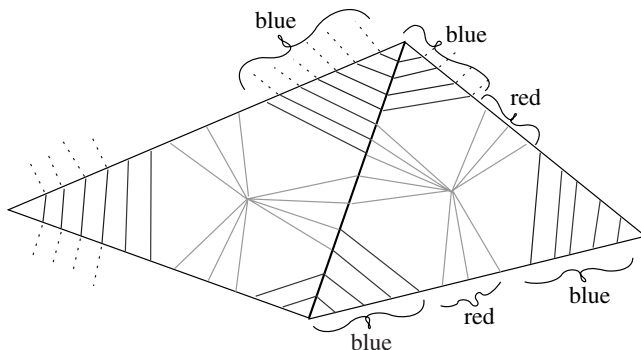


Figure 2. Two triangles in $\widetilde{P(G)}$, with blue and red segments. Consecutive blue segments are mapped into a single slice in $\mathcal{CE}(\Sigma)$, and red ones are mapped on segments of controlled length. Note that blue segments in a triangle can turn red in another.

To illustrate our construction, let us describe $\overline{\mathcal{BR}}$ in a triangle T of $P(G)$. The components of $T \setminus \overline{\mathcal{BR}}$ are either triangles with one blue edge (the spikes around the vertices of T , there are three of them, given T), or quadrilaterals with two parallel blue edges (there are arbitrarily many of those), or triangles with two red edges and one vertex being the singular red point of the triangle (there are at most $3 \times 10\psi(T(G))$ of them), or pentagons with two consecutive red edges (containing the red singular point) and the opposite edge blue (there are three of them, given T).

Lemma 3.2. *To every component C of $P(G)$ one may associate a graph dC with blue and red edges, having one or two connected components, so that the following holds:*

- (i) *In the disjoint union of all graphs dC , over all components C , there are at most $2T(G) \times 30\psi(T(G))$ red edges (twice the total number of red edges in $\overline{\mathcal{BR}}$).*
- (ii) *Let C be a component and $\partial C = \overline{C} \setminus C$. Any loop l_1 in dC is homotopic to a loop l_2 in ∂C such that l_1 has, for each color, at most twice the number of edges l_2 has.*
- (iii) *The embedding of a component of dC in C induces an epimorphism on the π_1*

except possibly if dC has one component, in which case $\pi_1(dC)$ maps in $\pi_1(C)$ either surjectively or in a subgroup of index 2.

(iv) The group $\pi_1(C)$ is free.

Proof. Let us consider C a connected component of $P(G) \setminus \overline{\mathcal{BR}}$. In C we choose dC to be the boundary of a small tubular neighborhood of $\partial C = \overline{C} \setminus C$. This is a graph, and for every triangle T of $P(G)$ and any component C_0 of $T \cap C$, $dC \cap C_0$ has same number of components as $\partial C \cap \overline{C_0}$ has (that is one or two). We color each of them by the color of the neighboring component of $\partial C \cap \overline{C_0}$, thus ensuring (i) since every edge of $\overline{\mathcal{BR}}$ locally separates $P(G)$ in two.

When homotoping dC to ∂C in the relevant tubular neighborhood of ∂C , we send edges of a given color on edges of ∂C of the same color and, again because every edge of $\overline{\mathcal{BR}}$ locally separates $P(G)$ in two components, send at most two edges of dC on the same edge of ∂C , hence (ii).

Let now p be a base point in dC and ℓ a loop in C starting at p . Let us denote by $T_1, T_2, \dots, T_k = T_1$ the consecutive triangles in which ℓ enters. For $i \leq k - 1$, we can inductively homotope ℓ so that it stays in the same component of $dC \cap T_i$. Then in T_k , either ℓ enters in the component of p or enters in the other component (if any). In the first case, ℓ is homotopic to a loop of dC , and in the second case ℓ^2 is homotopic to a loop of dC ; moreover in this second case, dC (globally) has only one connected component.

Thus, in the first case the inclusion induces an epimorphism on the fundamental groups. In the second case we need to show that if $\pi_1(dC)$ is not surjective in $\pi_1(C)$, then it is contained in a subgroup of index 2. The image of $\pi_1(dC)$ contains the subgroup S generated by all the squares of $\pi_1(C)$. This subgroup S is normal since a conjugate of a product of squares is a product of squares, and the quotient $\pi_1(C)/S$ is finitely generated (as $\pi_1(C)$ is) and has all its elements of order 2. Hence it is abelian, hence finite, and is isomorphic to $(\mathbb{Z}/2\mathbb{Z})^n$ for some n . Now, if $\pi_1(dC)$ is not surjective in $\pi_1(C)$, since it contains S , $\pi_1(dC)/S$ maps on some proper subgroup of $\pi_1(C)/S$. Since the quotient is abelian, one can form the quotient by this proper subgroup, which gives a certain $(\mathbb{Z}/2\mathbb{Z})^k$, $k \leq n$. Hence there is a surjective map on $\mathbb{Z}/2\mathbb{Z}$ that contains $\pi_1(dC)$ in its kernel. This establishes point (iii).

To obtain (iv), it suffices to notice that C is homotopically equivalent to $C' = C \setminus \{C_i\}$, where the C_i are the components of $C \cap T$, for some T , that are triangles. Now C' is a union of quadrilaterals and pentagons glued together on two opposite sides. Consider the graph with one edge in each such quadrilateral or pentagon joining the midpoints of the opposite sides on which the gluing is done. It is easily checked that C' is then homeomorphic to a fibration of an open interval $(-1, 1)$ on this graph, which makes its fundamental group free, which proves (iv). \square

Lemma 3.3. *Let X be the bipartite graph such that the vertices of one color (white) are the components of $\overline{\mathcal{BR}}$, the vertices of the other color (black) are the components of $P(G) \setminus \overline{\mathcal{BR}}$, and that the edges realize the adjacency relation.*

Then G has a natural structure of fundamental group of a finite graph of groups such that the underlying graph is X , the vertex groups are the fundamental groups of the relevant components, and the edge groups are the fundamental subgroups of the relevant intersections, with identifications, through edges, induced by the adjacency in $P(G)$.

Remark. The graph X is endowed with groups for each vertex and each edge, and attaching maps from edge groups to adjacent vertex group. Here one should note that the attaching maps need not be injective, and that the vertex and edge groups need not embed in G . This is not important if one is interested in finding a presentation of G using this construction. If one wants to find a graph of groups with all maps injective, one should take for vertex groups, instead of the fundamental groups of the components, their *images* in G .

Proof. This is an application of the van Kampen theorem to our decomposition of $P(G)$. \square

Let us now describe a particular generating set for the vertex and edge groups of this graph of groups.

Let c be a component of $\overline{\mathcal{BR}}$ (it is a red-and-blue graph in $P(G)$). We make a careful choice of generators (this construction actually works for any red-and-blue graph).

Let \mathcal{B} be the maximal blue subgraph of c (not necessarily connected). Let b_1, \dots, b_k be a minimal collection of blue edges (possibly empty) so that $\mathcal{B} \setminus \{b_1, \dots, b_k\}$ is a forest. By minimality of k , $\mathcal{B} \setminus \{b_1, \dots, b_k\}$ has same number of connected components as \mathcal{B} , hence $c \setminus \{b_1, \dots, b_k\}$ is connected. Let now r_1, \dots, r_n be a minimal collection (possibly empty) of red edges so that $c \setminus \{b_1, \dots, b_k, r_1, \dots, r_n\}$ is a tree. Such a collection exists, since if a red-and-blue graph is not a tree, it has a cycle that cannot consist only of blue edges if the maximal blue subgraph is a forest. Hence this red edge can be removed, and the new graph is still connected (and has less edges).

Note that $c \setminus \{b_1, \dots, b_k, r_1, \dots, r_n\}$ is in fact a maximal subtree of c , for, if we put back one red edge, it is not a tree, and if we put back one blue edge, some blue component is not a tree. Therefore, if $*$ is a base point in c , $\pi_1(c, *)$ has a natural isomorphism with the free group on $\{b_1, \dots, b_k, r_1, \dots, r_n\}$. We call the r_i the *red generators* and the b_j the *blue ones*.

Let now C be a component of $P(G) \setminus \overline{\mathcal{BR}}$. Each component of dC (see Lemma 3.2) is a red-and-blue graph, so the construction above can be performed, thus providing a system of generators of dC .

Lemma 3.4. *If the image of G in \mathcal{MCE} is purely pseudo-Anosov, then a loop defining a blue generator (in the graph dC of a component C or in a component of $\overline{\mathcal{BR}}$) is trivial in G .*

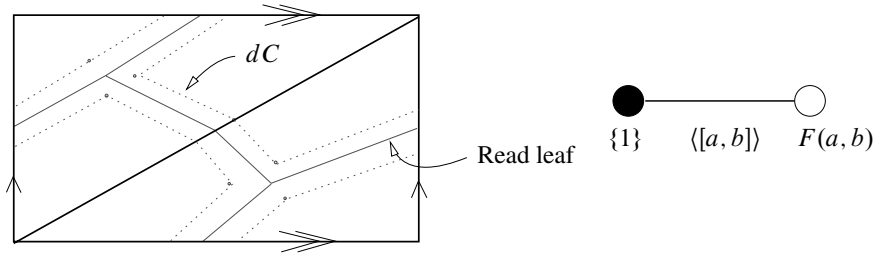


Figure 3. In the torus (left) with only one red leaf (surrounded by its tubular neighborhood defining dC), the graph X has a white vertex (corresponding to the red graph) and a black vertex (the 2-cell attached to it).

Proof. Let γ be a closed blue curve in $\overline{\mathcal{BR}}$ and note that $\varphi(\gamma)$ is defined up to conjugacy. By Lemma 3.1, $\varphi(\gamma)$ has a finite orbit in $\mathcal{CG}(\Sigma)$ (contained in the slice associated to the blue graph we started from). Thus $\varphi(\gamma)$ is not pseudo-Anosov, and by assumption this means that γ is trivial in G .

It remains to check that a loop defining a blue generator is freely homotopic to a blue curve in $P(G)$. Consider a generator associated to a blue edge e , let $B(e)$ be the blue graph containing e , and $T_b(e)$ the chosen maximal subtree of $B(e)$ (so $T_b(e)$ does not contain e). Let T be the chosen maximal subtree of the red-and-blue graph (it contains $T_b(e)$ but not e). By definition of the generator, its loop is a concatenation of three paths: a path p in T from the base point to e , the blue edge e itself, and finally a path q in T from e to the base point. We claim that p and \bar{q} (with reverse orientation) enter $T_b(e)$ on the same point. If it were not the case, a path p_b in $T_b(e)$ between the two entering points would give rise to a loop $(p)(p_b)(q)$ in T , contradicting that T is a tree. This implies also that from the base point to this entering point, the paths p and \bar{q} are equal. Hence, the loop is freely homotopic to a loop contained in the graph $B(e)$ which is completely blue. \square

In the following $b_1(G, \mathbb{Z}/2\mathbb{Z})$ is the first Betti number of G over $\mathbb{Z}/2\mathbb{Z}$. Let N_0 be the number of red edges in $\overline{\mathcal{BR}}$ in $P(G)$. Note that $N_0 \leq T(G) \times 30\psi(T(G))$ (see Remark after Lemma 3.1).

Lemma 3.5. *G has a presentation as a bipartite graph of groups X' (with black and white vertices as above) that satisfies the following properties.*

- *There are at most $T(G)$ white vertices whose groups are free with at most N_0 generators each.*
- *There are at most $2N_0$ black vertices whose groups are free with at most $2N_0$ generators each.*
- *For each edge and each adjacent vertex, the attaching map from the edge group into the vertex group sends each generator to a product of at most $2N_0$ of the given generators.*

- For each edge the attaching map from the edge group into the adjacent black vertex group is either surjective or has its image in a subgroup of index 2.
- There are at most $b_1(G, \mathbb{Z}/2\mathbb{Z})$ edges for which the attaching map into the adjacent black vertex group is not surjective.

Proof. The graph of groups X' is not necessarily the X defined above: first we simply remove every vertex of X with no red generator (and their adjacent edges), since by Lemma 3.4 their groups are trivial. Then, if an edge e of X has only blue generators (hence with trivial group), it is separating (otherwise G would have a cyclic free factor, hence several ends), and one of the components of $X \setminus e$ has trivial group (otherwise G has several ends). By removing this subgraph, we can assume that no edge has only blue generators without changing the fundamental group. Thus we get the graph of groups X' , and only now it is possible to bound the number of white vertices and black vertices, respectively, by the total number of components of $\overline{\mathcal{BR}}$ with a red edge (less than the total number of red singular points, $T(G)$) and by twice the number of red edges in $\overline{\mathcal{BR}}$ (each component of $P(G) \setminus \overline{\mathcal{BR}}$ associated to a vertex of X' is adjacent to a certain red edge and only two of them can be adjacent to the same red edge).

The generators of each vertex or edge groups are chosen to be the red ones in the construction above (the blue ones being all trivial by Lemma 3.4). The obtained graph then satisfies the two first points by construction.

Given an edge, it corresponds to an adjacency of a component C of $P(G) \setminus \overline{\mathcal{BR}}$ (its black vertex) and a component of $\overline{\mathcal{BR}}$ (its white vertex). Hence it corresponds to a component of dC .

By Lemma 3.2 (ii), each loop in a component of dC is homotopic to a loop in the relevant component of $\overline{\mathcal{BR}}$ containing at most twice the number of edges of each color. If the loop in dC is simple, it is homotopic to a loop of $\overline{\mathcal{BR}}$ passing at most twice through each edge it contains. Thus it is homotopic to the product of $2N_0$ red generators of the relevant component of $\overline{\mathcal{BR}}$, which proves the third point for attaching maps into white vertex groups. For attaching maps into black vertex groups, the bound is obtained similarly, replacing the component of $\overline{\mathcal{BR}}$ by the graph in C to which C is homotopically equivalent.

The fourth point follows by Lemma 3.2 (iii).

It remains to bound the number of edges whose group is in a subgroup of index 2 in the adjacent black vertex group: each of them gives rise to a morphism of G onto $\mathbb{Z}/2\mathbb{Z}$ (by sending the index 2 subgroup containing the edge group on 0, and its non-trivial coset in the white vertex on 1). One cannot have more than $b_1(G, \mathbb{Z}/2\mathbb{Z})$ distinct such morphisms. \square

In the graph X' we choose a base white vertex v_0 , and in $P(G)$ a base point p_0 in the component of $\overline{\mathcal{BR}}$ of v_0 . For each black vertex w of T' (which has valence at most 2), we choose a path between its two adjacent components of $\overline{\mathcal{BR}}$ as follows. By adjacency, there is a triangle of the initial triangulation of $P(G)$ in which these

components have adjacent segments. We choose a segment s_w between two such points in that triangle, so that s_w does not intersect any other component.

We thus can choose a simple path from p_0 to any component of $\overline{\mathcal{BR}}$ by a sequence of simple paths in components and segments s_w for black vertices w of X' . Once chosen a maximal subtree T' in X' , this gives a choice of one simple path from p_0 to any component, hence a choice of well-defined conjugacy classes of any vertex group of X' in G .

Thus, given a component c of $\overline{\mathcal{BR}}$, its base point p is chosen to be the first point of the component met by the former path, and the red generators of $\pi_1(c, p)$ are seen as elements of G .

Finally, we choose in the universal cover $\widetilde{P(G)}$ a pre-image \tilde{p}_0 of p_0 , and from it, cross sections of each of the chosen paths from p_0 to a component of $\overline{\mathcal{BR}}$. If p is a base point of a component, this gives \tilde{p} a particular pre-image of p in $\widetilde{P(G)}$.

The next lemma should be compared with [7], Lemma III.4.

Lemma 3.6. *Let c be a component of $\overline{\mathcal{BR}}$ that contains a red edge, and p the base point in c . Let $v \in \mathcal{CG}(\Sigma)$ be $v = \tilde{\varphi}(\tilde{p})$ with the notation just introduced.*

Then the image under φ of each of the red generators of $\pi_1(c)$ translates v by at most $T(G)(20\psi(T(G)) \times 1000\delta + 200\delta)$ in $\mathcal{CG}(\Sigma)$.

Moreover, if c' is another component that is adjacent to c in T' (their white vertices in T' are at distance 2), then the vertex v' for c' is at distance at most $1000\delta + T(G)(20\psi(T(G)) \times 1000\delta + 200\delta)$ from v in $\mathcal{CG}(\Sigma)$.

Proof. Let us choose from the point \tilde{p} a connected cross section \tilde{c} of c in $\widetilde{P(G)}$. This is a graph with blue and red edges, and $\tilde{\varphi}$ maps it in $\mathcal{CG}(\Sigma)$. Any red generator of $\pi_1(c)$ moves the base point of \tilde{c} to a point of \tilde{c} . The equivariance of φ implies that the image in $\mathcal{MCG}(\Sigma)$ of the generators of $\pi_1(c)$ move $v = \tilde{\varphi}(\tilde{p})$ (which we choose as the vertex v of the statement) to a point of $\tilde{\varphi}(\tilde{c})$.

Each blue segment of it is mapped in a finite subgraph (a slice in fact), with universally bounded diameter (at most 200δ , by Lemma 2.14 (ii)).

Each pair of red edges around a singular point (corresponding to the center of the slice decomposition of a triangle) is mapped on a path between to points at uniformly bounded distance (at most $20\psi(T(G)) \times 1000\delta$, by Proposition 2.15 and Lemma 2.14 (iii)).

Now in any segment of \tilde{c} , there are at most $T(G)$ different singular points, since it is a cross section of a graph on $P(G)$. Thus there are at most $T(G)$ pairs of red edges as above. Therefore, the extremal points of a segment in \tilde{c} are mapped by $\tilde{\varphi}$ to points at distance at most $T(G)(20\psi(T(G)) \times 1000\delta + 200\delta)$ from each other. This proves the first claim.

Now if c' is a component adjacent to c in T' , its base point is mapped in a slice that is adjacent to a slice of a point of $\tilde{\varphi}(\tilde{c})$, thus at distance at most 1000δ from this point by Lemma 2.14 (iii). With the former estimate on the diameter of $\tilde{\varphi}(\tilde{c})$, this gives the required bound. □

We can now prove Theorem 1.2.

Proof of Theorem 1.2. Let $D_0 = T(G) \times (20\psi(T(G)) \times 1000\delta + 200\delta)$ and $\text{Diam}(X')$ be the diameter of the graph X' ; the two constants depend only on G and Σ . The generating set is that given by the graph of the group X' in Lemma 3.5, taking red generators for every vertex and edge groups. From Lemma 3.5, we can write a presentation over this generating set: the relations are the words $s^{-1}f(s)$ where s runs over the generating sets of the edge groups, and $f(s)$ is the image of s under the corresponding attaching map (note that $f(s)$ is equal to a product of at most $2N_0$ generators of the range of this map).

Hence we have a bound on the complexity of the presentation over this generating set. From Lemma 3.6 we get a subset Δ of $\mathcal{CG}(\Sigma)$ of diameter bounded by $\text{Diam}(X') \times (D_0 + 1000\delta)$ such that for all generators r in our family, there is $v \in \Delta$ with $d(v, \varphi(r)v)$ universally bounded (by D_0 if the generator is in a vertex group, and by $\text{Diam}(X') \times (D_0 + 1000\delta)$ if the generator is a stable letter of the graph X' with the maximal subtree T'). The triangle inequality easily implies that for any point v in Δ and any generator r in our family, the displacement is bounded by $2 \text{Diam}(\Delta) + \text{Diam}(X') \times (D_0 + 1000\delta)$, which depends only on G and Σ . \square

References

- [1] J. Barnard, Bounding surface action on hyperbolic spaces. Preprint 2007. [arXiv:0709.2722](https://arxiv.org/abs/0709.2722)
- [2] B. H. Bowditch, Intersection numbers and the hyperbolicity of the curve complex. *J. Reine Angew. Math.* **598** (2006), 105–129. [Zbl 1119.32006](https://zbmath.org/journals/JRAM/598/105.html) [MR 2270568](https://mr.computer.org/doi/abs/10.1090/S0025-5718-2006-01720568)
- [3] B. H. Bowditch, Tight geodesics in the curve complex. *Invent. Math.* **171** (2008), 281–300. [Zbl 05248088](https://zbmath.org/journals/InvMat/171/281.html) [MR 2367021](https://mr.computer.org/doi/abs/10.1007/s00208-008-0201-1)
- [4] B. H. Bowditch, Atoroidal surface bundles over surfaces. Preprint 2007. <http://www.warwick.ac.uk/masgak/papers/bhb-atoroidal.pdf>
- [5] B. H. Bowditch, One-ended subgroups of mapping class groups. Preprint 2007. <http://www.warwick.ac.uk/masgak/papers/bhb-oneended.pdf>
- [6] F. Dahmani, Accidental parabolics and relatively hyperbolic groups. *Israel J. Math.* **153** (2006), 93–127. [Zbl 05256318](https://zbmath.org/journals/IJM/153/93.html) [MR 2254639](https://mr.computer.org/doi/abs/10.1007/s12220-006-9039-9)
- [7] T. Delzant, L'image d'un groupe dans un groupe hyperbolique. *Comment. Math. Helv.* **70** (1995), 267–284. [Zbl 0834.20038](https://zbmath.org/journals/CMH/70/267.html) [MR 1324630](https://mr.computer.org/doi/abs/10.1007/S0003680000030)
- [8] H. A. Masur and Y. N. Minsky, Geometry of the complex of curves I: Hyperbolicity. *Invent. Math.* **138** (1999), 103–149. [Zbl 0941.32012](https://zbmath.org/journals/InvMat/138/103.html) [MR 1714338](https://mr.computer.org/doi/abs/10.1007/S0020800900038)
- [9] E. Rips and Z. Sela, Canonical representatives and equations in hyperbolic groups. *Invent. Math.* **120** (1995), 489–512. [Zbl 0845.57002](https://zbmath.org/journals/InvMat/120/489.html) [MR 1334482](https://mr.computer.org/doi/abs/10.1007/S0020800900038)

Received March 13, 2008; revised February 2, 2009

F. Dahmani, Institut de Mathématiques de Toulouse, Université Paul Sabatier,
31062 Toulouse, Cedex 9, France

E-mail: francois.dahmani@math.univ-toulouse.fr

K. Fujiwara, Graduate School of Information Science, Tohoku University, Sendai,
980-8579, Japan

E-mail: fujiwara@math.is.tohoku.ac.jp