Relative hyperbolization of (one-ended hyperbolic)-by-cyclic groups

François Gautero and Martin Lustig

Abstract

We show that the semi-direct product of a one-ended torsion-free word-hyperbolic group with $\mathbb{Z}$, given through an automorphism $\alpha : G \to G$, is a hyperbolic group relative to certain canonical subgroups of G on which α acts periodically or with linear growth.

Introduction

Let G be a word-hyperbolic group in the sense of Gromov [17], and let α be an automorphism of G. For simplicity we consider in this paper only the case of torsion-free G. The overall question adressed in this paper is whether (or better, to what extend) the mapping torus group of α, i.e. the semi-direct product

$$G_\alpha = G \rtimes_\alpha \mathbb{Z},$$

is itself a word-hyperbolic group. The general answer to this question is mixed, pending on the properties of the automorphism α. The best way to phrase this is in terms of *relative hyperbolicity* (as introduced by Farb), explained in detail in section 6.

Theorem 0.1. *For any freely indecomposable non-elementary torsion-free word-hyperbolic group G and every automorphism α of G the semi-direct product $G \rtimes_\alpha \mathbb{Z}$ is word-hyperbolic relative to a family of subgroups which consist entirely of elements on which α acts up to conjugacy periodically or with linear growth.*

Notice that the statement of Theorem 0.1 is sharp in the sense that it becomes wrong whenever one tries to decrease the relative part essentially: Indeed, every element of an α-periodic conjugacy class in G is contained in a $\mathbb{Z} \oplus \mathbb{Z}$ subgroup of G_α. A similar statement is true for a linearily growing element, relative to the family of periodic subgroups. For any $\mathbb{Z} \oplus \mathbb{Z}$ subgroup, one of its summands must lie (up to conjugacy) in the relative part of G in order to get a relatively hyperbolic group G_α. These subgroups also explain why above we use Farb's version of relative hyperbolicity: In order for G_α to be relatively hyperbolic in the stronger sense of Gromov, both factors of any $\mathbb{Z} \oplus \mathbb{Z}$ subgroup must be contained in the relative part. Such a coarser "quotientation" of G_α, with a stronger relative hyperbolicity property, will be discussed in Remark 7.4.

Remark 0.2. If G is a non-trivial free product, then its automorphisms behave, relative to the free factors, very much like automorphisms of free groups (compare [7]), so that the results of [13] can be adapted. In combination with the results and methods of the present paper, with a small amount of extra technical work, [13] gives the prospective of a complete solution of our overall question for all automorphisms α of all (torsion-free) hyperbolic groups G, stated as follows: $G_\alpha = G \rtimes_\alpha \mathbb{Z}$ is word-hyperbolic relative to a canonical set of subgroups of G that consist entirely of (and contain up to conjugation all) elements that have polynomial α-growth.

Using Sela's decomposition of one-ended hyperbolic groups, described in section 7 below, the above Theorem 0.1 will follow directly from the more general statement of Theorem 7.1, where G is replaced by the fundamental group of a graph of groups of a certain type.

An important special case, which is at the time a "warm up" as well as the most important building block in the proof of these results, is that of a surface group G, where α is given by a surface homeomorphism.

Theorem 0.3. *Let M_h be the mapping cylinder of a homeomorphism $h\colon S \to S$ of a surface S with Euler characterisitc $\chi(S) < 0$, possibly with non-empty boundary ∂S. Then the fundamental group $\pi_1 M_h$ is word-hyperbolic relative to the family of subgroups defined (up to conjugation) by the fundamental groups of the boundary curves of S, of the reduction curves in the Nielsen-Thurston reduction of h, and of those subsurfaces of S which are maximal unions of components in this reduction on which h induces a periodic mapping class, pasted together along a maximal subfamily of the reduction curves.*

Before attacking the case of reducible h in section 6 we very carefully consider the "absolute" case of a closed surface S provided with a pseudo-Anosov homeomorphism $h : S \to S$. (Note that, if the mapping class of h is reducible or of finite order, then the mapping-torus M_h of h contains a π_1-injective 2-torus, and hence its fundamental group contains a copy of $\mathbb{Z} \oplus \mathbb{Z}$ and thus can not be word-hyperbolic.) The first five sections of this paper are devoted to a detailed analysis of the geometry of mapping tori of such pseudo-Anosov surface homeomorphisms. We assemble the tools necessary for the more difficult relative case, and, as a side product, we also obtain a conceptually simple proof of the following:

Proposition 0.4. *If the mapping class of a surface homeomorphism $h\colon S \to S$ is pseudo-Anosov, then $\pi_1 M_h$ is word-hyperbolic.*

Note that Proposition 0.4 is also a well known consequence of each of three important, highly non-trivial theories: of Thurston's hyperbolisation of fibered 3-manifolds [19], of Bestvina-Feighn's combination theorem [1], or of the work of Gabai-Kazez [12],[11] on genuine essential laminations in 3-manifolds.

By contrast, the methods presented in this paper only use classical Nielsen-Thurston surface theory [10] and elementary geometric considerations: We use the stable and the unstable singular foliation on S, which is associated to the pseudo-Anosov map h, to define a *foliation metric* on the universal cover of S. This foliation metric is lifted and properly rescaled to define a metric on the universal cover $\tilde{M}$ of M_h (reminiscent of the metric used by Cannon-Thurston, see [5, 8]). This metric is preserved by the free, properly discontinuous and cocompact action of $\pi_1 M_h$. Hence, in order to prove that $\pi_1 M_h$ is hyperbolic, it suffices to show that geodesic triangles in $\tilde{M}$ are thin (see [4] and [15] for background). This is done through a close investigation of our metric on $\tilde{M}$, which results in an approximation of geodesics by *vertical* and *horizontal* segments. We find this geodesic approximation illuminating in its own way, also with respect to potential further applications.

1 Setting

Throughout the end of section 5 we consider a compact, connected surface S with empty boundary and with Euler characteristic $\chi(S) < 0$, as well as a pseudo-Anosov homeomorphism h of S. The surface S can be endowed with a hyperbolic structure, given by an identification of the universal cover of S with the Poincaré disk D^2. The action of the fundamental group $\pi_1 S$ on the universal cover of S by deck transformations is realized via this identification by isometries on D^2 with respect to Poincaré's hyperbolic metric.

We choose a lift $\tilde{h}$ of h to D^2. It admits a pair of $\pi_1 S$-equivariant, $\tilde{h}$-invariant, transverse, singular foliations $\mathcal{F}_u$, $\mathcal{F}_s$ in D^2, called the *unstable* and the *stable* foliation respectively. Both $\mathcal{F}_u$ and $\mathcal{F}_s$ come with $\pi_1 S$-equivariant, transverse measures, denoted μ_u and μ_s. The transverse measure μ_s on $\mathcal{F}_s$ is uniformely dilated by a factor $\lambda > 1$ under the induced action of $\tilde{h}$, and similarly μ_u is uniformely contracted by the factor $\frac{1}{\lambda}$. In particular this gives for any finite segment l_0 of a leaf l of $\mathcal{F}_s$, and for any finite segment l'_0 of a leaf l' of $\mathcal{F}_u$, the equations (compare [10]):

$$\mu_u(\tilde{h}(l_0)) = \frac{1}{\lambda}\mu_u(l_0)\,, \qquad \mu_s(\tilde{h}(l'_0)) = \lambda\mu_s(l'_0)$$

The universal cover of the mapping torus $M_h = S\times[0,1]/\{(s,1) = (h(s),0)\mid s \in S\}$ is called *mapping telescope* and denoted by $\tilde{M}$. It is homeomorphic to $D^2 \times \mathbb{R}$, but for our purposes we have to carefully select the identification $\tilde{M} = D^2 \times \mathbb{R}$ among the different (natural) possibilities. This is most conveniently done by specifying a free, properly discontinuous and cocompact action of $\pi_1 M_h$ on $D^2\times\mathbb{R}$, and by identifying M_h with the quotient space:

Every element $g\tau^r$ of $\pi_1 M_h = \pi_1 S \rtimes_{\tilde{h}_*} \mathbb{Z}$, for $g \in \pi_1 S$, τ a generator of the infinite cyclic "right factor" $\mathbb{Z}$ and r any integer, can be identified with the lift $H = g\tilde{h}^r\colon D^2 \to D^2$ of the r-th power of the given homeomorphism $h\colon S \to S$. We define the corresponding "deck transformation" $\hat{H}$ on $\tilde{M}$ by

$$\hat{H}(g\tau^r) = \hat{H}\colon D^2 \times \mathbb{R} \to D^2 \times \mathbb{R},\ (x,t) \mapsto (H(x), t-r).$$

For $r = 0$ this gives the (desired) consequence that an element $g \in \pi_1 S$ acts on $\tilde{M} = D^2\times\mathbb{R}$ through $g(x,t) = (gx,t)$. In particular, the quotient manifold $(D^2 \times$

$\mathbb{R})/\pi_1 M_h$ is obtained in two steps as follows: First one quotients $D^2 \times \mathbb{R}$ modulo $\pi_1 S$ to obtain $S \times \mathbb{R}$, on which the infinite cyclic group $\pi_1 M_h/\pi_1 S$ generated by the image $\bar{\tau}$ of τ acts through $\bar{\tau}(z,t) = (h(z), t-1)$. As second step one quotients $S \times \mathbb{R}$ modulo $< \bar{\tau} >$, which clearly gives the mapping torus M_h.

As a consequence, for each *time* t the *fiber* $F_t = D^2 \times \{t\}$ of $\tilde{M}$ covers a fiber homeomorphic to S of the mapping torus fibration $M_h \to S^1$. More specifically, the covering map $D^2 \times \mathbb{R} \to M_h$ induces for every integer value $t \in \mathbb{Z} \subset \mathbb{R}$ a covering map $D^2 \times \{t\} \to S$, $(x,t) \mapsto y$ which differs from the above specified universal covering map $D^2 \to S$, $x \mapsto z$ by composition with the power h^t, i.e. $y = h^{-t}(z)$. The analogous statement is true for any two fibers F_{t_1}, F_{t_2} at times $t_1, t_2 = t_1 + t$ which differ by an integer $t \in \mathbb{Z}$.

We now rescale the transverse measures on the fibers continuously, by multiplying μ_u with the factor λ^{-t} to get the transverse measure of the unstable foliation on $D^2 \times \{t\}$, and by multiplying μ_s with factor λ^t to get the transverse measure of the unstable foliation on $D^2 \times \{t\}$. This has the effect that any of the above defined deck transformations $\hat{H}\colon D^2 \times \mathbb{R} \to D^2 \times \mathbb{R}$ preserves both transverse measures of any path which is contained in some fiber $F_t \subset \tilde{M}$.

We define the *telescope flow* $\Phi_t\colon \tilde{M} \to \tilde{M}$ by $\Phi_t(x,r) = (x, r+t)$. This is nothing other than the lift under the covering map of the mapping torus flow on M_h. If $P = (x,t)$ is a point in $\tilde{M}$, we will denote by $O_P = \{x\} \times \mathbb{R}$ its orbit under the telescope flow. Note that the above rescaling of the metric has the effect that two orbits O_P, O_Q in the same stable leaf will approach one another arbitrarily close and exponentially fast in the future (i.e. for $t > 0$), and they will diverge exponentially fast in the past. The converse is true for unstable leaves.

2 The foliation metric on the Poincaré disk

In the Poincaré disk D^2, identified above with the universal cover of the surface S, we use the two foliations $\mathcal{F}_u$, $\mathcal{F}_s$ and their transverse measures μ_u, μ_s to define a new metric:

The *unstable length* $|\gamma|_u$ (resp. *stable length* $|\gamma|_s$) of a path γ in D^2 is the total variation of $|d\mu_u|$ (resp. $|d\mu_s|$) along γ. The *foliation distance* $d_{\mathcal{F}}(x,y)$ between any two points x, y in D^2 is the infimum, over all the paths γ between x and y, of the *foliation length* $|\gamma|_{\mathcal{F}} = |\gamma|_u + |\gamma|_s$.

We denote by $D^2_{\mathcal{F}}$ the Poincaré disk equipped with the foliation distance $d_{\mathcal{F}}$. The latter is clearly equivariant with respect to the covering space action of $\pi_1 S$. As this action is free, properly discontinuous and cocompact, the disk $D^2_{\mathcal{F}}$ is quasi-isometric to the group $\pi_1 M_h$. In particular one has:

Proposition 2.1. *The disk $D^2_{\mathcal{F}}$ is a Gromov hyperbolic metric space.* □

Similarly, each fiber $F_t = D^2 \times \{t\}$ of $\tilde{M}$ can be equipped with the analogous *foliation metric*, denoted below by $d^t_{\mathcal{F}}$. The unstable and the stable foliation length in F_t will be denoted by $|\gamma|^t_u$ and $|\gamma|^t_s$ respectively. Since F_{t+1} is isometric to F_t (through the deck transformation $\hat{H}$ associated in section 1 to any lift H of the homeomorphism $h\colon S \to S$), and since the Gromov hyperbolicity constant δ_t for F_t varies continuously with $t \in \mathbb{R}$, there is an upper bound $\delta > 0$ to all δ_t.

The following is an elementary exercise (which becomes almost trivial, if one passes from measured foliations to the associated geodesic laminations).

Lemma 2.2. *(i) A path $\gamma\colon [0, d] \to D^2_{\mathcal{F}}$ is a geodesic if and only if it intersects any leaf of $\mathcal{F}_u \cup \mathcal{F}_s$ in such a fashion that the full γ-preimage of the intersection is connected.*

(ii) A geodesic γ intersects each leaf of $\mathcal{F}_u \cup \mathcal{F}_s$ which separates the endpoints of γ, and no other leaf.

(iii) Any two geodesics in $D^2_{\mathcal{F}}$ with same endpoints have the same unstable and the same stable length. □

Lemma 2.3. *Let l be a leaf of $\mathcal{F}_u$ which separates two points $A, B \in D^2_{\mathcal{F}}$, and let l' be a similar leaf of $\mathcal{F}_s$. If l and l' intersect, then the intersection point $Y \in D^2_{\mathcal{F}}$ is contained in some geodesic between A and B. In particular, any two geodesics $[A, Y]$ and $[Y, B]$ intersect only in Y, and their union is a geodesic $[A, B]$.*

Proof. Assume Y_0 is a point on both $[A, Y]$ and $[Y, B]$, and let l_0 be the leaf of $\mathcal{F}_u$ and l'_0 the leaf of $\mathcal{F}_s$ that contain Y_0. If $Y_0 \neq Y$, then by Lemma 2.2 (ii) l_0 or l'_0 separates Y from both A and B. But then one of the leaves l or l' that contain Y must also be separated by l_0 or l'_0 from both A and B. Hence one of them does not separate A from B, which contradicts the assumption on l or that on l'. □

Proposition 2.4. *Let $A, B, C \in D^2_{\mathcal{F}}$ be arbitrary three points. Then there exists a point $Z \in D^2_{\mathcal{F}}$ which lies on three properly chosen geodesics $[A, B]$, $[B, C]$ and $[A, C]$. In particular there is a geodesic triangle $\Delta = [A, B, C]$ which is a tripod with center Z.*

Proof. Any leaf of either $\mathcal{F}_u$ or $\mathcal{F}_s$ which intersects one side of a geodesic triangle $[A, B, C]$ must also intersect at least one other side. One observes that there must be precisely one *special* leaf in either foliation which intersects all three sides of $[A, B, C]$. The special leaf can be regular, in which case it must contain at least one vertex of the triangle. According to the location of the intersection points of the two special leaves from $\mathcal{F}_u$ and $\mathcal{F}_s$ with the sides of $[A, B, C]$ one distinguishes several cases, and for each of them one sees directly that the two special leaves must intersect. Hence we can apply Lemma 2.3 to derive that the intersection point Z of the two special leaves has all the desired properties. □

From the arguments in the last proof it follows also that the center Z is uniquely determined by the three given points A, B, C, but this fact is not used here.

3 The telescope distance on $\tilde{M}$

For any two points $P, Q \in \tilde{M}$ and for any time $t \in \mathbb{R}$ let $P_t \in O_P$ and $Q_t \in O_Q$ be the intersection points of the fiber F_t with the orbits O_P and O_Q respectively.

From the definition of the foliation metric in F_t in the last section it follows directly that, unless P_t and Q_t belong to a leaf of $\mathcal{F}_s$ or of $\mathcal{F}_u$, the distance $d^t_{\mathcal{F}}(P_t, Q_t)$ becomes minimal at the *diagonal time* t_{diag} *for* P *and* Q, by which we mean the time coordinate $t = t_{diag}$ of the fiber F_t where the stable and the unstable distances between P_t and Q_t are equal. This distance is called the *orbit width* of O_P and O_Q and denoted by $\omega(P, Q)$. The points $P_{t_{diag}}$ and $Q_{t_{diag}}$ are said *to lie in diagonal position.* Away from the diagonal fiber $F_{t_{diag}}$the distance $d^t_{\mathcal{F}}(P_t, Q_t)$ increases with t into both directions, in an exponential (or more precisely, a hyperbolic cosin like) fashion.

If P_t and Q_t are contained in the same leaf of $\mathcal{F}_s$ or of $\mathcal{F}_u$, we set $\omega(P, Q) = 0$ and $t_{diag} = \infty$ or $t_{diag} = -\infty$.

We say that two points P', R' in some fiber F_t are *in* (ϵ, δ)*-almost diagonal position*, for $\epsilon \geq 0$, if there are points P and R in the δ-neighborhood of P' and R' respectively, which are in diagonal position and which satisfy $d^t_{\mathcal{F}}(P, R) \geq \epsilon$. We observe:

Lemma 3.1. *For any sufficiently large $\epsilon > 2\delta$ there exists a constant $\sigma > 0$ such that for any two points P', R' in some fiber F_t which are in (ϵ, δ)-almost diagonal position, the diagonal time t' for P', R' satisfies $|t - t'| \leq \sigma$.* □

A path γ in $\tilde{M}$ is called *vertizontal* if it is the concatenation of finitely many subpaths that are *vertical* (i.e. segments of some orbit O_P) or *horizontal* (i.e. contained in some fiber F_t). The length $l(\gamma)$ of such a path is simply the sum of the length of their subpaths, where each horizontal subpath $\gamma' \subset F_t$ is measured in the foliation metric $d^t_{\mathcal{F}}$ defined on F_t, while the length of a vertical segment is simply the difference of the time coordinates of its endpoints.

Lemma 3.2. *Let γ be a vertizontal path in $\tilde{M}$ from $P = (x_P, t_P)$ to $R = (x_R, t_R)$, and assume that γ realizes at the point $Q_{min} = (x_{min}, t_{min})$ the minimal time coordinate. Let γ_{min} be the projection of γ under the telescope flow to the fiber $F_{t_{min}}$. Then the length of γ satisfies the inequality:*

$$l(\gamma) \geq |\gamma_{min}|^{t_{min}}_u + t_P + t_R - 2t_{min}$$

Similarly, if $Q_{max} = (x_{max}, t_{max})$ is the point of γ with the maximal time coordinate, then one has, for the projection $\gamma_{max} \subset F_{t_{max}}$ of γ:

$$l(\gamma) \geq |\gamma_{max}|^{t_{max}}_s - t_P - t_R + 2t_{max}$$

Proof. Between P and Q_{min} the path γ accumulates a time contribution of at least $t_P - t_{min}$, and between Q_{min} and R it accumulates similarly at least $t_R - t_{min}$. From the scaling of the foliation metric in the fibers according to time we obtain that the unstable contribution to $l(\gamma)$ is bounded below by $|\gamma_{min}|^{t_{min}}_u$, which proves the first inequality. The second one is completely analogous. □

We define a metric d on $\tilde{M}$ (the *telescope metric*) as follows:

Definition 3.3. *For any two points $P, Q \in \tilde{M}$ let $d(P, Q)$ denote the infimum of the lengths $l(\gamma)$ of all vertizontal paths γ joining P to Q.*

Using a decomposition of S into rectangles with sides given by stable and unstable leaves (as done in the usual Markov partitions for $h\colon S \to S$, see [10]) it is not hard to see that this infimum is indeed realized by some vertizontal path, which hence is a geodesic $[P, Q]$. Thus $\tilde{M}$ is a proper geodesic metric space on which $\pi_1 M_h$ acts by isometries in a free, cocompact and properly discontinuous way, so that $\tilde{M}$ and $\pi_1 M_h$ are quasi-isometric. In particular, in order to show that $\pi_1 M_h$ is word-hyperbolic, it suffices to show that geodesic triangles in $\tilde{M}$ are thin.

Lemma 3.4. *The image $\gamma_t \subset F_t$ of a geodesic $\gamma \subset \tilde{M}$ under the telescope flow is a geodesic with respect to the foliation metric $d^t_{\mathcal{F}}$.* $\square$

Proof. If γ_t is not a geodesic, then by Lemma 2.2 (ii) it intersects some leaf l' of the stable or the unstable foliation which does not separate the endpoints of γ_t. In particular, γ_t intersects twice another leaf l that separates the endpoints of γ_t from l'. Hence we can strictly reduce the length of γ by "shortcutting" a subpath γ' with endpoints in the full lift $L \subset \tilde{M}$ of l, and γ' intersects the full lift of l': We replace γ' with a path γ'' that is entirely contained in L. Here γ'' is the projection of γ' into L along leaves of the stable or the unstable foliations in the fibers F_t which intersect γ'. $\square$

A path $\gamma = [A, B] \subset \tilde{M}$ from A to B is called a *geodesic lift* (of γ_t), if the telescope flow projection γ_t of γ into any fiber F_t is a geodesic in F_t with respect to the metric $d^t_{\mathcal{F}}$, and if there is no strictly shorter path in $\tilde{M}$ from A to B which also projects to γ_t. Clearly every geodesic in $\tilde{M}$ is a geodesic lift. Similarly as shown above for geodesics one can see that a vertizontal geodesic lift exists for any $d^t_{\mathcal{F}}$-geodesic in a fiber F_t.

4 Geodesics in $\tilde{M}$

The goal of this section is to show that any geodesic in $\tilde{M}$ locally either has an "almost vertical" direction, in which case it stays close to an orbit O_X, or else, if it has a "mostly horizontal" direction, it passes close to the diagonal projection of a properly chosen subsegment of the geodesic. The precise statement is given by the following three propositions.

For any $c > 0$ we say that a subset S of $\tilde{M}$ is a *c-approximation* of a subset $T \subset \tilde{M}$, if S and T are contained in the c-neighborhood of each other (i.e. if their Hausdorff distance is smaller or equal to c). We say that a point $Q \in \tilde{M}$ is contained in the *horizontal c-neighborhood* of an orbit O_P if in the fiber F_Q that contains Q there is a path γ_Q which connects Q to $P' = O_P \cap F_Q$ and is of length $l(\gamma_Q) \leq c$.

Proposition 4.1. *For any $d > 0$ there exists a bound $c = c(d) > 0$ with the following property, for any points $X, Y \in \tilde{M}$:*

If Y is contained in the horizontal d-neighborhood of the orbit O_X, then any geodesic lift $\gamma = [X, Y]$ is contained in the horizontal c-neighborhood of O_X.

Furthermore, γ is c-approximated by the vertical segment $[X, X'] \subset O_X$, where $X' \in O_X$ is the point in the same fiber F_Y as Y.

Proof. For any point $Z \in \gamma$ let t_Z denote the time coordinate of Z. We consider the projection γ_Z of the initial segment $[X, Z]$ of γ into the fiber F_Z that contains Z. It suffices to show that t_Z lies in a neighborhood of $[t, t_Y] \subset \mathbb{R}$ which depends only on d, and that both the stable and the unstable length of γ_Z are bounded above by a constant which also depends only on d.

We consider the vertizontal path γ' which is the union of the vertical segment $[X, X']$ with the projection $\gamma_Y \subset F_Y$ of γ, and observe:

$$l(\gamma) \leq l(\gamma') \leq |t_X - t_Y| + d$$

By symmetry we can assume that $t_X \geq t_Y$, so that the above estimate shows that γ must be contained between the fibers $F_{t_Y - \frac{d}{2}}$ and $F_{t_X + \frac{d}{2}}$. This establishes the desired bounds for t_Z. We also deduce, for any point $Z \in \gamma$, that the stable length of γ_Z is bounded above by the stable length of the projection $\gamma_{t_Y - \frac{d}{2}}$ of γ to $F_{t_Y - \frac{d}{2}}$ which depends only on d:

$$|\gamma_Z|_s^{t_Z} \leq |\gamma_{t_Y - \frac{d}{2}}|_s^{t_Y - \frac{d}{2}} \leq \lambda^{\frac{d}{2}} d$$

To get the bound on the unstable length of γ_Z we apply Lemma 3.2 with $P = Z$ and $R = X$ and obtain, for the smallest time coordinate t_{min} of any point $Q_{min} \in [X, Z]$, and for the projection γ_{min} of $[X, Z]$ to the fiber $F_{t_{min}}$, the inequality

$$l([X, Z]) \geq |\gamma_{min}|_u^{t_{min}} + t_X + t_Z - 2t_{min}$$

and thus

$$(t_Z - t_Y) + (|\gamma_{min}|_u^{t_{min}} + t_X + t_Z - 2t_{min}) \leq l([Y, Z]) + l([Z, X]) =$$

$$l(\gamma) \leq l(\gamma') \leq t_X - t_Y + d\,,$$

which gives

$$|\gamma_{min}|_u^{t_{min}} + 2(t_Z - t_{min}) \leq d\,.$$

From the last inequality we deduce directly the desired upper bound

$$|\gamma_Z|_u^{t_Z} = \lambda^{t_Z - t_{min}} |\gamma_{min}|_u^{t_{min}} \leq \lambda^{\frac{d}{2}} d$$

which depends only on d. □

Lemma 4.2. *(a) For any $h > 0$ there exists $j = j(h) > 0$ with the following property: Let $\gamma = [P, R]$ be any geodesic lift in $\tilde{M}$ that joins orbits of width $\omega(P, R) \leq h$, and let F_t be the fiber closest to P with the property that the projection $\gamma_t \subset F_t$ of γ has length $l(\gamma_t) \leq h$. Then γ passes through the j-neighborhood of γ_t.*
(b) For any $0 < h_0 \leq h_1$ there exists a bound $j(h_0, h_1) > 0$ such that for any h with $h_0 \leq h \leq h_1$ one has $j(h) \leq j(h_0, h_1)$.

Proof. (a) If the endpoints of γ lie on different sides of F_t, then the claim is obvious, as there is a point where γ has to meet F_t and hence γ_t. Hence we can assume that γ lies completely on one side of F_t, and by symmetry we can assume that $t < t_{min}$, where we use the same notation as in Lemma 3.2. The latter gives $l(\gamma) \geq |\gamma_{min}|_u^{t_{min}} + t_P + t_R - 2t_{min}$. Note that the assumption $t < t_{min}$ implies also that $t \geq t_{diag}$, where t_{diag} denotes the time coordinate of the diagonal projection of γ, so that we deduce that the unstable length of γ_t is bigger than the stable one and hence also bigger than $\frac{h}{2}$.

Now we consider the vertizontal path γ' which descends from the endpoint P of γ along the orbit O_P down to F_t, then proceeds along the projection γ_t, and finally reconnects to the other endpoint R of γ along the orbit O_R. The length of γ' is equal to $l(\gamma') = t_P - t + t_R - t + l(\gamma_t) \leq t_P - t + t_R - t + h$, which has to be greater or equal to the length of the geodesic lift γ. Hence we obtain $|\gamma_{min}|_u^{t_{min}} + t_P + t_R - 2t_{min} \leq l(\gamma) \leq l(\gamma') \leq t_P - t + t_R - t + h$, thus $2(t_{min} - t) + h \geq |\gamma_{min}|_u^{t_{min}} \geq \lambda^{t_{min}-t}|\gamma_t|_u^t \geq \lambda^{t_{min}-t}\frac{h}{2}$, and finally

$$4(t_{min} - t) + (2 - \lambda^{t_{min}-t})h \geq 0$$

which implies that $t_{min} - t$ is bounded above by some $j > 0$ which only depends on h.
(b) The statement (b) is a direct consequence of the last inequality. □

Notice that in the previous lemma the fiber F_t can not be replaced by the diagonal fiber for the orbits O_P and O_R: In fact, for small values of h the distance between these two fibers can become arbitrary large. Note also that F_t depends on P and on O_R, but not on the location of R on O_R. The same is true in part (ii) of the following:

Proposition 4.3. *(i) For any $h > 0$ there exists $k = k(h) > 0$, such that any geodesic lift $\gamma = [P, R]$ in $\tilde{M}$ with orbit width $\omega(P, R) < h$ is contained in the horizontal k-neighborhood of the orbits $O_P \cup O_R$.*
(ii) Furthermore, γ is k-approximated by the union of two vertical segments $[P, P_t] \subset O_P$ and $[R, R_t] \subset O_R$, where the fiber F_t that contains P_t and R_t is the fiber closest to, say, P where the projection $\gamma_t \subset F_t$ of γ has length $l(\gamma_t) \leq h$.

Proof. We apply Lemma 4.2 (a) to obtain a bound $j = j(k)$ a point $Q \in \gamma$ which lies in the j-neighborhood of the fiber F_t described in part (ii) above. The projection γ_t of γ to the fiber F_t has length $l(\gamma_t) \leq h$, so that the projection γ_Q of γ to the fiber F_Q that contains Q has length $l(\gamma_Q) \leq \lambda^j h$.

Thus we can apply Proposition 4.1, for $d = \lambda^j h$, to both subsegments $[P, Q]$ and $[R, Q]$ of γ and obtain as direct consequence the desired statements (i) and (ii), for $k = c(\lambda^j h)$. □

Proposition 4.4. *For any sufficiently large $k > 0$ there exists a bound $m = m(k) > 0$, such that, for any geodesic lift $\gamma = [P, R]$ in $\tilde{M}$ and any point $Q \in [P, R]$ which lies on the boundary of the horizontal k-neighborhoods of both the orbit O_P and the orbit O_R, the distance from Q to the diagonal projection γ_{diag} of γ is smaller or equal to m.*

Proof. Fix $h > 0$. From Proposition 4.3 (i) we know that, if $\omega(P, R) < h$, then Q is contained in the horizontal $k(h)$-neighborhood of $O_P \cup O_R$. Thus, for $k > k(h)$ we conclude from the stated hypotheses that $l(\gamma_{diag}) = \omega(P, R) \geq h$, while the same hypotheses give $l(\gamma_Q) = 2k$ for the horizontal projection γ_Q of γ which contains Q. But since orbits in distinct stable and unstable leaves diverge exponentially in the future and in the past, see section 3, there is an upper bound $m = m(k)$ to the distance between $Q \in \gamma_Q$ and γ_{diag}. □

5 Geodesic triangles in $\tilde{M}$ are thin

Proposition 5.1. *There exists a constant $b > 0$ such any two geodesic lifts γ, γ' in $\tilde{M}$ with same endpoints A, B are b-approximations of each other.*

Proof. Fix a constant $\epsilon > 2\delta$ as in Lemma 3.1, let $k(\epsilon)$ be the constant provided by Proposition 4.3, choose $k' > k(\epsilon)$, and define $d' = 3k'$.

If the orbit width $\omega(A, B)$ is smaller than d', then both γ and γ' are $k(d')$-approximated by the same vertical subsegments of O_A and O_B, by Proposition 4.3, and hence they are $2k(d')$-close.

If the orbit width satisfies $\omega(A, B) \geq d'$, then, since $k' = \frac{d'}{3}$, there is a point $Q \in \gamma$ outside of the horizontal k'-neighborhood of the endpoint orbits O_A and O_B. Thus there exist points $P, R \in \gamma$ such that $Q \in [P, R] \subset \gamma$, and Q lies on the boundary of the horizontal k'-neighborhoods of both orbits O_P and O_R. Thus Proposition 4.4 proves that Q lies in the $m(k')$-neighborhood of the diagonal projection $[P_{diag}, R_{diag}]$ at time t_{diag} of the subsegment $[P, R]$ of γ. Note also that the geodesic segment $[P, R]$ of γ projects to a segment $[P_t, R_t]$ in the fiber F_t containing Q which has length $d^t_{\mathcal{F}}(P_t, R_t) = 2k'$, so that we deduce $\omega(P, R) \leq 2k'$.

Our choice of the constant $k' > k(\epsilon)$ allows us to deduce from Proposition 4.3 (i) that $\omega(P, R) \geq \epsilon$. Furthermore, we know from section 2 that the fiber F_{diag} which contains $[P_{diag}, R_{diag}]$ is δ-hyperbolic, with respect to the foliation metric $d^{t_{diag}}_{\mathcal{F}}$. Hence the geodesic γ' contains points $\hat{P}, \hat{R}$ which are projected by the telescope flow to points $P', R' \in F_{diag}$ in the horizontal δ-neigborhood of P_{diag} and R_{diag} respectively. Lemma 3.1 proves that the diagonal time t' for P' and R' satisfies $|t_{diag} - t'| \leq \sigma$. This gives the possibility to determine upper and lower bounds for $\omega(\hat{P}, \hat{R})$ as follows: $2k' + 2\delta \geq \omega(P, R) + 2\delta \geq d^{t_{diag}}_{\mathcal{F}}(P', R') \geq \omega(\hat{P}, \hat{R}) \geq \lambda^{-\sigma} d^{t_{diag}}_{\mathcal{F}}(P', R') \geq \lambda^{-\sigma}(\omega(P, R) - 2\delta) \geq \lambda^{-\sigma}(\epsilon - 2\delta)$. We then apply Lemma 4.2 (b) with $h_1 = 2k' + 2\delta$ and $h_0 = \lambda^{-\sigma}(\epsilon - 2\delta)$ and $h = \omega(\hat{P}, \hat{R})$. The last equation implies that in Lemma 4.2 (a) the fiber F_t agrees with the diagonal fiber $F_{t'}$ for $\hat{P}$ and $\hat{R}$. Thus there is a point $Q' \in \gamma'$ of distance $\leq j(h_0, h_1)$ from the diagonal projection of the segment $[\hat{P}, \hat{R}]$, which itself has distance $\leq \sigma$ from $[P_{diag}, R_{diag}]$, which in turn has length $\leq 2k'$ and contains a point of distance $\leq m(k')$ from Q. This gives the desired uniform upper bound on $d(Q, Q')$.

Still in the case $\omega(A, B) \geq d'$, for any $P \in \gamma$ in the horizontal k'-neighborhood of the orbit O_A (or of O_B) there is a point Q as above on the boundary of this horizontal neighborhood, with $P \in [A, Q]$. Then Proposition 4.1 shows that P lies within horizontal distance $c = c(k')$ from a point on the vertical segment $[A, A_t]$ of O_A, where t is the time coordinate of Q. Now, the point $Q' \in \gamma'$ constructed above is of bounded distance from Q, so that Proposition 4.1 provides a vertical approximation $[A, A']$ of $[A, Q']$ with A' not far from A_t. Hence P is close to some point P' of γ'. □

Proposition 5.2. *There exists a constant $k > 0$ such that every geodesically lifted triangle $\Delta = [A, B, C]$ in $\tilde{M}$, which projects to a tripod in any fiber, is k-thin.*

Proof. We proceed in close analogy to the proof (though not to the statement) of Proposition 5.1: Outside a horizontal neighborhood of the orbit of the triangle vertices and of the tripod center O_Z we conclude as in the above considered case of geodesically lifted digons. Inside these orbit neighborhoods we use Propositions 4.1 (in the case of large orbit width between O_Z and the vertex orbit) and 4.3 (in the case of small orbit width) to approximate the segments of the sides of Δ by vertical segments on that orbit. The endpoints of these orbit segments are determined by the vertices of Δ or by the time coordinate of a point Q on the boundary of the above horizontal neighborhood, if such a point exists (i.e. in the case of large orbit width, where Proposition 4.1 is applied). If for one of the vertices of Δ, say A, the orbit width $\omega(A, Z)$ is small, then the corresponding endpoints of the orbit segments on O_A and on O_Z are determined by the time coordinate of $A_t \in O_A$ given as the points $P_t \in O_P$ in Proposition 4.3 (ii). Notice here that we crucially exploit the fact that the fiber F_t in Proposition 4.3 (ii) depends on P and on O_R but not on the particular location of R on the orbit O_R, as a priory we do not know whether the two sides $[A, B]$ and $[A, C]$ of Δ meet the orbit O_Z at a close distance.

Thus an endpoint of such a vertical segment is either independent from the side of Δ which is being approximated, or else it is closeby to the endpoint of the vertical segment, on the same orbit, which approximates one of the other two sides of Δ. In any case each of the vertical segments is contained in a bounded neighborhood of the corresponding vertical segment from the other side or from the other two sides of Δ. □

Theorem 5.3. *There exists a constant $h > 0$ such that every geodesic triangle $\Delta = [A, B, C]$ in $\tilde{M}$ is h-thin.* □

Proof. From Proposition 2.4 we know that there is another triangle Δ' with same endpoints as Δ, but with sides that are only geodesic lifts, such that Δ' is projected by the telescope flow to a tripod, in any fiber F_t. Now Proposition 5.2 shows that Δ' is thin, while Proposition 5.1 shows that Δ' is an approximation of Δ, so that Δ must also be thin. □

Theorem 5.3 shows that $\tilde{M}$ is hyperbolic in Gromov's sense. As the action of $\pi_1 M_h$ on $\tilde{M}$ is isometric, free, properly discontinuous and cocompact (see section 3), this implies Proposition 0.4 from the introduction as an immediate corollary.

Remark 5.4. An interesting feature of the material presented in this and the previous sections comes from the fact that foliations with a transverse measure determine a dual $\mathbb{R}$-tree. Our proof of Theorem 5.3 generalizes nicely to this more general setting, so that direct extensions of this approach seem natural. This concerns for example certain classes of free group automorphisms, where however a main problem remains to be solved: The action of F_n on the product of two $\mathbb{R}$-trees $T_+ \times T_-$ is, in general, not cocompact and hence one has to find a suitable F_n-invariant subspace of $T_+ \times T_-$. Recently V. Guirardel has obtained interesting results on this matter, see [18].

6 Relative hyperbolic mapping tori: The special case

The goal of this and the next section is to prove Theorem 0.1 from the introduction, or rather, the stronger Theorem 7.1. For this purpose we need to extend the technique (and the result) of the previous sections to the case of reducible surface homeomorphisms, replacing "hyperbolicity" of the mapping torus by "relative hyperbolicity". There are two different notions of relative hyperbolicity in the literature, and they are note quite equivalent, though very close. We will give more details below in Remark 7.4. Until then we will use "relative hyperbolic" always in the sense of Farb's paper [9] (called "weakly hyperbolic" in [3]), which we recall now.

Let G be a finitely generated group, and let $\mathcal{U} = (U_1, \dots, U_r)$ be a finite family of subgroups U_i of G. A *Cayley graph* Γ_* *of* G *relative to* $\mathcal{U}$ is the Cayley graph Γ of G with respect to a finite generating system of G, together with extra vertices added, one for each left coset gU_i, and with extra edges added that connect any such extra vertex to any vertex of Γ which is labelled by an element of this coset.

One says that G is *relative hyperbolic with respect to* $\mathcal{U}$ if a Cayley graph of G relative to $\mathcal{U}$ is a δ-hyperbolic space. It is immediate that any two Cayley graphs of G relative to the same family of subgroups $\mathcal{U}$ are quasi-isometric, so that relative hyperbolicity depends only on G and $\mathcal{U}$. Similarly, replacing any U_i of $\mathcal{U}$ by a conjugate gU_ig^{-1} or by a finite index subgroup $U_i' \subset U_i$ changes the relative Cayley graph only up to a quasi-isometry.

The canonical action of G on a relative Cayley graph Γ_* is not any more free, as is the G-action on "absolute" Cayley graphs. However, every G-orbit in Γ_* is discrete, and on all but finitely many of them G acts freely. The exceptional orbits Gx_i are in 1:1 relation with the subgroups U_i of $\mathcal{U}$, in that $\mathrm{Stab}_G(x_i) = g_iU_ig_i^{-1}$ for some suitable $g_i \in G$. In particular, one easily verifies that the action of G on Γ_* satisfies the hypotheses of the following proposition, which we will use below. It seems to be (within an ϵ-neighborhood of) folk knowledge among the experts and can be proved easily by a proper variation of the well known proof in the classical case of empty $\mathcal{U}$, see e.g. [16]:

Proposition 6.1. *Let X be a geodesic metric space, and let G be a finitely generated group which acts on X by isometries. Assume that X/G can be covered by finitely many connected sets B_i with pairwise connected intersections, and with the property that some (and hence every) connected component $\tilde{B}_i$ of the full preimage of B_i in X has finite diameter and is convex. Assume furthermore that there is a lower bound $\epsilon > 0$ to the distance between disjoint lifts $g\tilde{B}_i$ and $h\tilde{B}_j$ for any $g, h \in G$. Let $U_i \subset G$ be the stabilizer of $\tilde{B}_i$.*

Then X is quasi-isometric to the group G relative the family $\mathcal{U}$ of subgroups U_i (i.e. to any Cayley graph of G relative to $\mathcal{U}$). □

In order to concentrate on the main difficulty in translating our techniques of the previous sections from absolute to relative hyperbolicity we will first consider a special case. In the next section, however, we shall see that this special case is also the main building block for the general relative hyperbolic case.

The special case considered in this section is that of a compact surface S_* of negative Euler characteristic with non-empty boundary ∂S_*. Let h_* be a pseudo-Anosov homeomorphism of S_*. As in Section 1, S_* is given a hyperbolic structure, and consequently $\pi_1 S_*$ acts by isometries on the hyperbolic plane, represented by the Poincaré disc D^2. The action, however, is in the present case not cocompact: The (finitely many) boundary curves of S_* are represented by geodesics on the hyperbolic surface $D^2/\pi_1 S$, and without loss of generality we can identify S_* with the geodesic compact core of $D^2/\pi_1 S_*$, i.e. with the subsurface of $D^2/\pi_1 S_*$ obtained by cutting off the infinite-volume funnels along the closed simple geodesics that represent the homotopy class of the boundary curves in $D^2/\pi_1 S_*$. We denote by D^2_* the subdisk of D^2 which is the full preimage of $S_* \subset D^2/\pi_1 S_*$ under the covering map $D^2 \to D^2/\pi_1 S_*$.

As in the closed case, we obtain $\tilde{h}_*$-invariant foliations $\mathcal{F}_u, \mathcal{F}_s$ in D^2_*, with projectively invariant measures μ_u and μ_s. There is, however, a substantial difference to the closed case, in that the boundary curves of S_* lift to leaves which belong to both foliations $\mathcal{F}_u$ and $\mathcal{F}_s$. At all other points of D^2_* the two foliations are transverse as before.

We continue to proceed, nevertheless, just as in the closed case, thus simply ignoring the above "difficulty": We define the mapping telescope $\tilde{M}_*$ as universal covering space of the mapping torus M_{h_*} and identify it with $D^2_* \times \mathbb{R}$. The action of $\pi_1 M_{h_*}$ on $\tilde{M}_* = D^2_* \times \mathbb{R}$ is defined precisely as described in section 1 for the closed case. Similarly, we copy from section 1 the notion of *fibers* and *orbits* in $\tilde{M}_*$, and also the rescaling of the measure on the fibers so that the action of $\pi_1 M_{h_*}$ on $\tilde{M}_*$ preserves the transverse measure of any horizontal path in $\tilde{M}_*$.

Next we introduce on D^2_* the *foliation distance* as described at the beginning of section 2. Note that, because of the above "problem", rather than getting a metric, this defines a pseudo-metric space $D^2_{*\mathcal{F}}$, where any two distinct points have distance 0 if and only if they belong to the same boundary geodesic of D^2_*. Let $\widehat{D}^2_*$ denote the canonical metric quotient space of D^2_*. We prove:

Proposition 6.2. *The disk $\widehat{D^2_{*\mathcal{F}}}$ is a Gromov hyperbolic metric space.*

Proof. From Proposition 6.1 we know that $\widehat{D^2_{*\mathcal{F}}}$ is quasi-isometric to the free group $\pi_1 S_*$, relative to the set of cyclic subgroups generated by the boundary curves (which are well defined up to conjugacy). Cyclic subgroups are known to be always quasi-convex in any hyperbolic group. By [14] any hyperbolic group G is hyperbolic relative to any quasi-convex subgroup. If G is a free group (as is the fundamental group of a surface with boundary), this follows alternatively from elementary arguments, which also extend directly to the case of more than one relative subgroup, see [13]. $\square$

As in section 2 one then defines the foliation (pseudo-)metric on each fiber F_t of $\tilde{M}_*$, and Lemmas 2.2 and 2.3 as well as Proposition 2.4 carry over word by word. Indeed, the same is true for all of sections 3 – 5, so that we obtain:

Theorem 6.3. *There exists a constant $h > 0$ such that every geodesic triangle $\Delta = [A, B, C]$ in $\tilde{M}_*$ is h-thin.* $\square$

Thus we can again apply Proposition 6.1 to obtain:

Corollary 6.4. *The group $\pi_1 M_{h_*}$ is hyperbolic relative to the family of cyclic subgroups generated by the boundary components of the surface, i.e. of the fiber.*

The reader should note here that this result is not a special case of Farb's relative hyperbolicity result for hyperbolic 3-manifolds with cusps (see [9]), as in the above result the relative subgroups are all infinite cyclic, while Farb considers hyperbolicity relative all of the fundamental group of any cusp neighborhood, which is isomorphic to $\mathbb{Z} \oplus \mathbb{Z}$.

7 Relative hyperbolic mapping tori: The general case

Let $\mathcal{G}$ be a finite graph of groups with fundamental group G, and let α be an automorphism of G which respects the graph of groups structure of $\mathcal{G}$: There are different equivalent ways to express this, the most elegant saying that the length function $||\cdot||_{\mathcal{G}}$, defined by $\mathcal{G}$ on the conjugacy classes of G, is invariant under α. Equivalently, there is a homeomorphism $H : T \to T$ of the Bass-Serre tree T associated to $\mathcal{G}$ which satisfies $\alpha(g)H = Hg : T \to T$ for every $g \in G$. Such an automorphism α induces a graph isomorphism α_Γ on the graph $\Gamma = T/G$ on which $\mathcal{G}$ is built, and for every vertex v and every edge e of Γ group isomorphisms $\alpha_v : G_v \to G_{\alpha_\Gamma(v)}$ and $\alpha_e : G_e \to G_{\alpha_\Gamma(e)}$ such that for the terminal vertex $v = \tau(e)$ of any (oriented) edge e and the corresponding group monomorphism $f_e : G_e \to G_v$, given by $\mathcal{G}$, one has $f_{\alpha_\Gamma(e)}\alpha_e = \alpha_v f_e$ up to inner automorphisms of $G_{\alpha_\Gamma(v)}$. Supposing the existence of such maps α_Γ, α_v and α_e gives another equivalent way to describe that α preserves the graph of groups structure of $\mathcal{G}$, see e.g. [6].

Let now V_0 be an α_Γ-invariant subset of the vertices of Γ, such that each of the groups G_v for $v \in V_0$ is the fundamental group of a surface S_v with boundary. We assume furthermore that for all edges e with terminal vertex $\tau(e) \in V_0$ the edge group G_e maps under the edge monomorphism f_e into an infinite cyclic subgroup of $G_{\tau(e)}$ that is generated (up to conjugation) by some boundary curve $\partial_e S_{\tau(e)}$ of the surface $S_{\tau(e)}$. In particular, these edge groups G_e must be infinite cyclic. (One could also allow trivial G_e, at the expense of minor technical modifications below.)

Finally, for each vertex $v \in V_0$ we assume that there is a homeomorphism $h_v : S_v \to S_{\alpha_\Gamma(v)}$ that induces the group isomorphism $\alpha_v : G_v \to G_{\alpha_\Gamma(v)}$. We denote by t_v the cardinality of the α_Γ-orbit of v. We then want to assume that the t_v-th power of the set of homeomorphisms h_v defines on S_v a homeomorphism $h_v^{(t_v)} = h_{\alpha_\Gamma^{t_v-1}(v)} h_{\alpha_\Gamma^{t_v-2}(v)} \dots h_{\alpha_\Gamma(v)} h_v$ which is pseudo-Anosov. We call the vertices from V_0 vertices *of type pseudo-Anosov on surface with boundary.*

We now define *complementary subgroups* as follows: For every boundary component $\partial_i S_v$ of a surface S_v with $v \in V_0$ one introduces a new vertex $v(\partial_i S_v)$ with vertex group isomorphic to $\mathbb{Z}$. The new vertex is connected by a newly introduced edge $e(\partial_i S_v)$ to v. The new edge group $G_{e(\partial_i S_v)}$ is also isomorphic to $\mathbb{Z}$, and on one side of $e(\partial_i S_v)$ the edge group monomorphism surjects $G_{e(\partial_i S_v)}$ onto $G_{v(\partial_i S_v)}$, while on the other side $G_{e(\partial_i S_v)}$ is mapped onto the infinite cyclic group $\pi_1 \partial_i S_v$. This does not change $\pi_1 \mathcal{G}$, and now all "old" edges of $\mathcal{G}$ adjacent to S_v can be sled over one of the new edges to be attached to one of the new vertices, again without changing $\pi_1 \mathcal{G}$. We now remove all vertices from V_0 together with all new edges (but not the new vertices !), and define the subgroups *complementary to the vertex subset* V_0 to be the fundamental groups of the obtained connected components of the left over sub-graph of groups. Each component determines a subgroup of $\pi_1 \mathcal{G}$ that is well defined up to conjugation.

We can now formulate the main result of this paper, in the most general form:

Theorem 7.1. *Let G be the fundamental group of a finite graph of groups $\mathcal{G}$, and let α be an automorphism of G which respects the structure of $\mathcal{G}$. Let V_0 be the set of vertices of type pseudo-Anosov on surface with boundary. Then the semidirect product $G_\alpha = G \rtimes_\alpha \mathbb{Z}$ is a word-hyperbolic group relative to the family of subgroups complementary to V_0.*

A direct consequence of this theorem is Theorem 0.3 from the introduction, where the graph of groups structure of $G = \pi_1 S$ is given by the reduction curves of the Nielsen-Thurston decomposition of the mapping class of h.

In order to derive Theorem 0.1, we use Sela's canonical graph of groups decomposition of one-ended hyperbolic groups G, as described and specified in [2]. "Canonicity" of this decomposition implies that it is preserved by any automorphism α of G. There are three kind of vertices in this decomposition: For torsion-free G the vertex groups are either *rigid*, infinite cyclic, or *quadratically hanging.* The edge groups are all infinite cyclic and they contain the infinite cyclic vertex groups. (The latter are only of auxiliary nature, in that one could get rid of them by contracting suitably chosen edges, at the expence of loosing canonicity.) The quadratically hanging vertex groups are fundamendal groups of surfaces with boundary (and thus free groups), and the adjacent edge groups are given (up to passage to a finite index subgroup) by the boundary components of the surface. Note that for torsion free hyperbolic G the hypotheses "one-ended" and "freely indecomposable non-elementary" are equivalent, by Stalling's celebrated characterisation of infinitely ended groups as amalgamated products with finite amalgam [20].

Thus Theorem 7.1 gives directly:

Theorem 7.2. *(a) Let G be a freely indecomposable non-elementary torsion-free word-hyperbolic group, and let α be an automorphism of G. Assume that in the canonical graph of groups decomposition of G on all of the quadratically hanging vertex groups the restriction of (a power of) α is induced by a pseudo-Anosov homeomorphism. Then the semi-direct product $G \rtimes_\alpha \mathbb{Z}$ is word-hyperbolic relative to the family of subgroups complementary to the set of vertices with quadratically hanging vertex group.*
(b) If the above assumption on the quadratically hanging vertex groups is not true, then the analogous statement is however valid, if the canonical decomposition of G is refined to include as further edges the reduction curves in the Nielsen-Thurston reduction of the induced automorphisms on the quadratically hanging vertex groups. The relative subgroups are then the subgroups complementary to the set of vertices which correspond to those subsurfaces on which α is pseudo-Anosov. □

Rigid vertex groups can be characterized by the fact that they are non-elementary and have finite outer automorphism group. Thus a power of α fixes every conjugacy class of elements in the vertex groups of the relative part, i.e. of any connected component of the subgraph of groups complementary to the set V_0 of vertices of type pseudo-Anosov on surface with boundary, in the refined canonical graph of groups $\mathcal{G}$, after having properly introduced the "new" vertices and edges as described before Theorem 7.1. It follows that such an automorphism is a multiple Dehn twist in the sense of [6] on every subgroup of G complementary to the set of vertices from V_0. Such multiple Dehn twists are well known to have linear or zero growth on all conjugacy classes. Thus the above Theorem 7.2 (together with Proposition 0.4 to include the case where $\mathcal{G}$ consists of a single vertex) implies directly Theorem 0.1 from the introduction.

To prove Theorem 7.1 we construct a topological space X modelled on the graph Γ in the usual way, to get $G = \pi_1 X$. Here each vertex v from V_0 is represented by the surface S_v, while all other vertex spaces $X_{v'}$ can be taken arbitrary (with $\pi_1 X_{v'} = G_{v'}$, of course). The edge spaces X_e with $\tau(e) \in V_0$ are given by annuli which, in the forward direction of e, are attached to a power of the boundary curve $\partial_e S_v$ of the surface S_v. The universal covering $\tilde{X}$ of X, together with the standard action of the fundamental group $G = \pi_1 X$ by deck transformations, gives then a model for the Bass-Serre tree T associated to $\mathcal{G}$, where each vertex of T is represented by a connected component of the full preimage in $\tilde{X}$ of a vertex space, and conversely. This is quite standard though a little heavy on the notational side, as is typical for Bass-Serre theory.

We now provide each of the surfaces S_v with the h_v-invariant measured foliations and lift the latter to each connected component of the full preimage of S_v in $\tilde{X}$. We then continue for each of these connected components precisely as in the special case considered in the previous section, while all the connected components of the full preimage of any other vertex space or of any edge space is given the zero pseudo-metric, thus effectively contracting each such connecting component. Let $\widehat{X}$ denote the obtained metric space.

As in the special case treated in the previous section, we now pass to the mapping telescope $\widehat{M} = \widehat{X} \times \mathbb{R}$, with fibers and orbits as defined before, and with an action of $G \rtimes_\alpha \mathbb{Z}$ defined on it so that, with respect to the properly rescaled metric on the fibers and the induced telescope metric on $\widehat{M}$, this action is isometric. Hence, in light of Proposition 6.1, in order to prove Theorem 7.1 it suffices to show the following:

Proposition 7.3. *There exists a constant $h \geq 0$ such that all geodesic triangles in the metric space $\widehat{M}$ are h-thin.*

Proof. We use the tree like structure of $\widehat{M}$, inherited from T, to observe that a geodesic triangle is either contained in a vertex space of $\widehat{M}$, in which case we are done, by the special case considered in Theorem 6.3 in the last section. Otherwise, the telescope flow projection of the triangle in any fiber decomposes into finitely many geodesic digons, each contained in a single vertex space, and one geodesic triangle, also contained in a single vertex space. Hence the claim follows again directly from the special case of surfaces with boundary, if we can bound the time coordinate difference which the lifts of the geodesic digons realize on their two sides at the place when they pass from one vertex space to another, i.e. at a "singular orbit" which corresponds to a maximal connected union of edge space and vertex spaces that are not of type pseudo-Anosov on surface with boundary (i.e. a conncted component of subspaces that has been contracted above).

To obtain such a bound we first use Proposition 5.1 to reduce the claim to bigons with the same telescope flow projections in any fiber. One then uses precisely the same arguments as explained in the proof of Proposition 5.2 for the time coordinate difference of two different geodesic lifts at the tripod center orbit O_Z. For digons with endpoints that have small orbit width and which have initial points (i.e. the endpoints closer to the vertices of the given triangle) on the two digon sides already at different time coordinate values, the time coordinate difference at the terminal endpoints can not be larger: This follows from the observation that, in the situation of Proposition 4.3 (ii), moving the point P along its orbit by a distance $\leq t_0$, will effect the time coordinate of the fiber F_t by at most t_0, as follows directly from the given definition of F_t. $\square$

Remark 7.4. There is a stronger version of relative hyperbolicity in the literature, originally proposed by Gromov [17], and later elaborated by Szczepański [21] and Bowditch [3]. It is equivalent (see [21]) to Farb's version of relative hyperbolicity, as used above, plus an aditional property, called "Bounded Coset Penetration property (BCP)" (also due to Farb [9]). The latter, however, implies that the relative subgroups U_i must be malnormal in the total group G. In our case, this is in general not true, as typically U_i is a cyclic summand of a $\mathbb{Z} \oplus \mathbb{Z}$ subgroup in G.

Thus, to get the stronger version of relative hyperbolicity for hyperbolic-by-cyclic groups, one is forced to increase the relative subgroups, so that with each periodic-up-to-conjugacy element it must contain at least its centralizer, which contains a subgroup isomorphic to $\mathbb{Z} \oplus \mathbb{Z}$. This is done properly by defining new relative subgroups $\widehat{U}_i = U_i \rtimes_{\alpha^t|_{U_i}} \mathbb{Z}$, where α^t is a power of α that fixes every vertex or edge group of G (up to conjugation). It follows then indeed that the group G_α from Theorem 7.1 is strongly relative hyperbolic relative to the family of these subgroups $\widehat{U}_i$. To show this one uses again the canonical decomposition of word-hyperbolic groups, where the arguments become actually simpler than in the above treated amalgamation: As the time coordinate must be collapsed as well, there is no need to worry about at which height one enters or leaves a vertex of type pseudo-Anosov on surface with boundary.

To prove the strong relative hyperbolicity of such vertex groups (i.e. the analogue of Corollary 6.4), one has two alternatives: One can either use the telescope metric on the mapping telescope $\tilde{M}_*$ as introduced in the last section, modify it properly by contracting each boundary component of $\tilde{M}_*$, and then use Lemma 1 of [21] to approximate geodesics in the resulting metric by geodesics with respect to the unmodified telescope metric. One then uses Propositions 4.1 and 4.3 and the existence of a lower bound to the distance between any two distinct boundary leaves to prove Farb's property BCP. Alternatively, one can simply quote [9], thus of course using implicitely Thurston's hyperbolisation of fibered 3-manifolds with pseudo-Anosov monodromy [19].

References

[1] M. BESTVINA and M. FEIGHN, 'A combination theorem for negatively curved groups', *Journal of Differential Geometry* **35** (1992), 85–101.

[2] B. BOWDITCH, 'Cut points and canonical splittings of hyperbolic groups', *Acta Mathematica* **180** (1998), 145–186.

[3] B. BOWDITCH, 'Relatively hyperbolic groups', preprint.

[4] M. BRIDSON and A. HAEFLIGER, *Metric spaces of non-positive curvature*, Fundamental Principles of Mathematical Science **319**, Springer Verlag 1999.

[5] J. CANNON and W.P. THURSTON, 'Group Invariant Peano Curves', preprint.

[6] M. M. COHEN et M. LUSTIG, 'The conjugacy problem for Dehn twist automorphisms of free groups', *Comment. Math. Helv.* **74** (1999), 179–200.

[7] D.J. COLLINS and E.C. TURNER, 'Efficient representatives for automorphisms of free products', *Michigan Math. Journal* **41** (3) (1994), 443–464.

[8] D. COOPER, D.D. LONG and A.W. REID, 'Bundles and finite foliations', *Inventiones Mathematicae* **118** (1994), 255–283.

[9] B. FARB, 'Relatively hyperbolic groups' *Geom. Funct. Anal.* **8** (1998), 1–31.

[10] A. FATHI, F. LAUDENBACK and V. POENARU, *Travaux de Thurston sur les surfaces*, Astérique **66–67**, 1979.

[11] D. GABAI, 'The Ubiquitous Nature of Quasi-Minimal Semi-Euclidean Laminations in 3-manifolds', *Surveys in Differential Geometry 5*, International Press.

[12] D. GABAI and W.H. KAZEZ, 'Group negative curvature for 3-manifolds with genuine laminations', *Geometry and Topology* **2** (1998), 65–77.

[13] F. GAUTERO, 'Hyperbolicité relative des mapping-torus d'automorphismes des groupes libres', preliminary preprint.

[14] S. GERSTEN, 'Subgroups of word-hyperbolic groups in dimension 2', *Journal of the London Mathematical Society* **54** (2) (1996), 261–283.

[15] E. GHYS and P. DE LA HARPE, *Sur les groupes hyperboliques d'après Mikhael Gromov*, Progress in Mathematics **83**, Birkhauser, 1990.

[16] E. GHYS and P. DE LA HARPE, 'Infinite groups as geometric objects (after Gromov)', *Ergodic theory, symbolic dynamics, and hyperbolic spaces*, 299–314, Oxford Univ. Press, 1991.

[17] M. GROMOV, 'Hyperbolic groups', *Essays in Group Theory* Math. Sci. Res. Inst. Publ. 8, Springer 1987, 75–263.

[18] V. GUIRARDEL, 'Intersection space for group actions on trees', preliminary preprint.

[19] J.P. OTAL, *Le théorème d'hyperbolisation pour les variétés fibrées de dimension* 3, Astérique **235**, 1996.

[20] J. STALLINGS, 'On torsion-free groups with infinitely many ends', *Annals of Mathematics* **88** (2) (1968), 312–334.

[21] A. SZCZEPAŃSKI, 'Relatively Hyperbolic Groups', *Michigan Math. J.* **45** (1998), 611–618.

University of Geneva, CH-1211 Geneva 24, Switzerland
Email adress: francois.gautero@math.unige.ch

Mathématiques (LATP), Université d'Aix-Marseille III, Ave. Escad. Normandie-Niemen, 13397 Marseille 20, France
Email adress: martin.lustig@univ.u-3mrs.fr