

THE OUTER SPACE OF A FREE PRODUCT

VINCENT GUIARDEL, GILBERT LEVITT

ABSTRACT. We associate a contractible “outer space” to any free product of groups $G = G_1 * \cdots * G_q$. It equals Culler-Vogtmann space when G is free, McCullough-Miller space when no G_i is \mathbb{Z} . Our proof of contractibility (given when G is not free) is based on Skora’s idea of deforming morphisms between trees.

Using the action of $\text{Out}(G)$ on this space, we show that $\text{Out}(G)$ has finite virtual cohomological dimension, or is VFL (it has a finite index subgroup with a finite classifying space), if the groups G_i and $\text{Out}(G_i)$ have similar properties. We deduce that $\text{Out}(G)$ is VFL if G is a torsion-free hyperbolic group, or a limit group (finitely generated fully residually free group).

1. INTRODUCTION AND STATEMENT OF RESULTS

A famous theorem of Grushko implies that a finitely generated group G has a decomposition as a free product of subgroups $G = G_1 * \cdots * G_p * F_k$, with F_k free of rank k and every G_i non-trivial, non isomorphic to \mathbb{Z} , and freely indecomposable. If $G = H_1 * \cdots * H_q * F_\ell$ is another such decomposition, the number of factors is the same, $\ell = k$, and (after reordering) H_i is conjugate to G_i .

Despite this uniqueness, there is a lot of freedom in the choice of the free factors, even when $k = 0$ (but $p \geq 3$). Because of this, the automorphism group $\text{Aut}(G)$ of G is much more complicated than the direct product of the groups $\text{Aut}(G_i)$ and $\text{Aut}(F_k)$. In particular, $\text{Aut}(G)$ contains automorphisms acting on G_i as conjugation by an element of G_j , with $j \neq i$, and acting on the other factors as the identity. A presentation of $\text{Aut}(G)$ was given by Fouxe-Rabinovitch [11], generalizing work of Nielsen on $\text{Aut}(F_n)$.

A different approach was introduced by Culler-Vogtmann [9], who obtained finiteness results for $\text{Out}(F_n)$ by letting it act on a contractible complex now known as outer space (see [30]). McCullough-Miller [23] constructed a complex to study the group $\Sigma\text{Aut}(G)$ of symmetric automorphisms of a free product. This group equals $\text{Aut}(G)$ only when no factor is infinite cyclic, so in a sense these two works cover opposite situations.

In this paper we construct a complex \mathcal{PO} that allows studying $\text{Aut}(G)$ and $\text{Out}(G)$ in all cases. We prove for instance:

Theorem 5.2. *Let $G = G_1 * \cdots * G_p * F_k$ be a Grushko decomposition of a finitely generated group G . Assume that $\text{Out}(G)$ is virtually torsion-free. If G_i and $\text{Aut}(G_i)$ have finite virtual cohomological dimension (resp. have a finite index subgroup with a finite classifying space) then so does $\text{Out}(G)$.*

Since a torsion-free finitely generated group is a free product of cyclic groups and one-ended groups, knowledge of $\text{Out}(H)$ for H one-ended gives information about automorphisms of arbitrary torsion-free groups.

We pointed out earlier that, in spite of the uniqueness in Grushko's theorem, there is a lot of flexibility in free products. On the other hand, one-ended groups often exhibit a strong algebraic rigidity. For instance, Bowditch showed [4] how to obtain the JSJ splitting of a one-ended (word) hyperbolic group H from the topology of its boundary. This implies that the splitting is completely unique, not just up to certain moves, and therefore invariant under automorphisms. Using this invariance, one may describe $\text{Out}(H)$ in terms of abelian groups and mapping class groups (see [20] for a precise statement of this result of Sela's [27]).

We deduce:

Theorem 6.1. *If G is a torsion-free hyperbolic group, then $\text{Out}(G)$ has a finite index subgroup with a finite classifying space. If G is a virtually torsion-free hyperbolic group, then $\text{Out}(G)$ has finite virtual cohomological dimension.*

Similarly, a limit group (finitely generated fully residually free group) has an invariant JSJ splitting if it is one-ended, and one gets:

Theorem 6.5. *If G is a limit group, then $\text{Out}(G)$ has a finite index subgroup with a finite classifying space.*

Let us now explain how the complex \mathcal{PO} is constructed. Consider a Grushko decomposition $G = G_1 * \cdots * G_p * F_k$ as above. Like Culler-Vogtmann's outer space, \mathcal{PO} may be viewed as a space of G -trees, or as a space of marked metric graphs. We describe it here as a space of graphs.




FIGURE 1. A graph of groups in the outer space of $G_1 * G_2 * G_3 * F_3$

A point in \mathcal{PO} is defined by a marked graph of groups Γ as in fig 1. More precisely, Γ is a finite graph of groups with trivial edge groups; its edges are assigned a positive length; the marking is an isomorphism from G to $\pi_1(\Gamma)$, well defined up to composition with inner automorphisms; for each $i \in \{1, \dots, p\}$ there is a vertex v_i whose group is conjugate to G_i ; all other vertex groups are trivial and every terminal vertex is a v_i .

A point of $P\mathcal{O}$ is such a marked metric graph, projectivized (all edge lengths may be multiplied by the same amount). The group $\text{Out}(G)$ acts on $P\mathcal{O}$ by change of marking. The quotient consists of finitely many open cells.

As always, the difficulty is to show contractibility of the space. Our approach is not combinatorial as in [9] or [23], but geometric. It is based on the idea of deforming morphisms between trees, introduced in Skora's unpublished paper [29] giving a geometric proof that Culler-Vogtmann space is contractible. We shall give the construction in its natural context, that of \mathbb{R} -trees, but in this paper we only apply it to metric simplicial trees.

Let $f : T_0 \rightarrow T_\infty$ be a map between metric simplicial trees, sending edges isometrically onto edges. We factor it through intermediate trees T_t ($t \geq 0$) defined as follows: two points $x, y \in T_0$ with $f(x) = f(y)$ are identified in T_t if and only if the image of the segment $[x, y]$ by f is contained in the t -ball around $f(x)$. The map f factors through maps $f_t : T_t \rightarrow T_\infty$. (We are grateful to L. Mosher for making us realize that this definition is equivalent to Skora's.)

One may also visualize the trees T_t as follows. If f is not an embedding, then it folds two edges having a common vertex. Let l be the length of the shortest pair of edges of T_0 folded by f . Then, for $t \leq l$, the tree T_t is defined by folding along length t any pair of edges that are folded by f . For $t > l$ with $t - l$ small enough, one can similarly define T_t from T_l , using the map f_l and folding along length $t - l$. It may be shown that one can reach any value $t \in \mathbb{R}^+$ by iterating this process finitely many times.

This deformation is constructed and studied in Section 3 (after the present paper was first posted, we became aware of [7], which contains another account of Skora's idea; see [21] for yet another deformation). Its most delicate feature is continuity.

Theorem. *The assignment $(f, t) \mapsto f_t$ is a continuous semi-flow on the space of morphisms between \mathbb{R} -trees (equipped with the equivariant Gromov-Hausdorff topology).*

In Section 4, we define our outer space and we show:

Theorem. *For any finitely generated group $G = G_1 * \cdots * G_p * F_k$, with $p \geq 1$, the outer space $P\mathcal{O}$ is contractible.*

This is proved as follows. We find a basepoint $T_0 \in P\mathcal{O}$ such that, for all $T \in P\mathcal{O}$, one can define a morphism $f_T : T_0 \rightarrow T$ varying continuously with T . Applying the semiflow to f_T (backwards) then gives a continuous way of deforming T into T_0 (to be precise, we work in the non-projectivized space \mathcal{O} , and the space is contracted to a simplex rather than to a point, as the edge lengths on T_0 depend on T).

As pointed out in [23], there are two natural topologies on $P\mathcal{O}$. First, there is the weak topology associated to its natural simplicial structure. Unlike Culler-Vogtmann space, the complex $P\mathcal{O}$ is usually not locally finite, and the weak topology is different from the equivariant Gromov-Hausdorff topology, or axes topology.

The deformation argument gives contractibility in the axes topology, whereas applications require weak contractibility. This is obtained by showing that, as T_0 and T each vary within a given simplex, the set of intermediate trees T_t only meets finitely many simplices.

The main difference between our arguments and those used by Skora in his proof that Culler-Vogtmann space is contractible [29] is that constructing the maps f_T is easier in our situation, because selecting the unique point of T fixed by G_1 provides a continuous choice of a basepoint in T as T varies in \mathcal{O} .

Because of this, contractibility of Culler-Vogtmann space is not proved here. In [15], we prove contractibility of arbitrary deformation spaces, using a general basepoint argument (two simplicial G -trees are in the same deformation space if they have the same elliptic subgroups, see [10]). This applies in particular to Culler-Vogtmann space, and to spaces of JSJ splittings.

In Section 5, we obtain general finiteness results about $\text{Out}(G)$ by studying its action on $P\mathcal{O}$. In Section 6, this is applied to hyperbolic groups, limit groups, and groups acting freely on \mathbb{R}^n -trees.

2. TREES

Trees will be considered both as combinatorial objects, and as metric objects.

Combinatorially, a (non-metric) simplicial tree is a simplicial 1-complex with no circuit. It is a G -tree if the group G acts on it by automorphisms, without inversions. A map between simplicial trees is *simplicial* if it maps every edge onto an edge (in particular, no edge is collapsed).

An \mathbb{R} -tree may be defined as a connected metric space whose distance d satisfies the 0-hyperbolicity inequality

$$d(x, y) + d(z, t) \leq \max(d(x, z) + d(y, t), d(x, t) + d(y, z))$$

(see [1]). Actions on \mathbb{R} -trees will always be by isometries.

Let T be a simplicial G -tree with finitely many G -orbits of edges. It becomes an \mathbb{R} -tree when each edge is assigned a positive length. An \mathbb{R} -tree obtained in this way will be called a *metric simplicial tree*. A metric simplicial tree is thus determined by the underlying simplicial tree, and one positive number for each orbit of edges.

The intermediate trees T_t will be defined as \mathbb{R} -trees, but readers uncomfortable with \mathbb{R} -trees may assume all trees to be simplicial. In this paper, the semi-flow will only be used with T_0 and T_∞ simplicial. In this case, all intermediate trees T_t are simplicial.

The segment between two points a and b in a tree will be denoted by $[a, b]$. In an \mathbb{R} -tree, the length of a segment $J = [a, b]$ will be denoted by $|J|$. A *finite subtree* of a tree is the convex hull of a finite set of points.

Let T be a G -tree. The tree, or the action, is called *non-trivial* if there is no global fixed point, *minimal* if there is no proper invariant subtree. In a minimal simplicial tree, there are only finitely many orbits of edges if G is finitely generated.

Maps between G -trees will always be assumed to be G -equivariant. A map $f : T \rightarrow T'$ between \mathbb{R} -trees is a *morphism* if each segment in T can be written as a finite union of subsegments, each of which is mapped isometrically into T' by f . If T and T' are metric simplicial trees, f is a morphism if and only if one may subdivide T and T' so that f becomes simplicial and maps every edge isometrically.

We always assume morphisms $f : T \rightarrow T'$ to be surjective. This is automatically true when T' is a minimal G -tree.

3. THE SEMI-FLOW

3.1. Definition of the semi-flow.

Let $f : T_0 \rightarrow T_\infty$ be a morphism between \mathbb{R} -trees. Our goal is to construct for each $t \in \mathbb{R}^+$ an \mathbb{R} -tree T_t , and morphisms $\varphi_t : T_0 \rightarrow T_t$ and $\psi_t : T_t \rightarrow T_\infty$, such that the following diagram commutes:

$$\begin{array}{ccc} T_0 & \xrightarrow{f} & T_\infty \\ & \searrow \varphi_t & \nearrow \psi_t \\ & T_t & \end{array}$$

If f is an equivariant morphism between two G -trees, then T_t is a G -tree and φ_t, ψ_t are equivariant.

When we need to study how this construction depends on f and t , we will write $T_t(f)$ instead of T_t , and $\Phi(f, t), \Psi(f, t)$ instead of φ_t, ψ_t .

The tree T_t as a set.

Let d_0 (resp. d_∞) denote distance in T_0 (resp. T_∞). Given $a, b \in T_0$ with $f(a) = f(b)$, we define their identification time by

$$\tau(a, b) = \max_{x \in [a, b]} d_\infty(f(x), f(a)).$$

Note that $\tau(a, b) \leq \frac{d_0(a, b)}{2}$.

For $t \geq 0$, say that $a \sim_t b$ when $f(a) = f(b)$ and $\tau(a, b) \leq t$. Obviously $\tau(a, c) \leq \max(\tau(a, b), \tau(b, c))$ if $f(a) = f(b) = f(c)$, so that \sim_t is an equivalence relation. We define T_t as T_0/\sim_t and $\varphi_t : T_0 \rightarrow T_t$ as the quotient map. The map f induces a map $\psi_t : T_t \rightarrow T_\infty$. The definitions may be extended to $t = \infty$, by setting $\varphi_\infty = f$ and $\psi_\infty = id_{T_\infty}$.

The metric on T_t .

This metric will be the maximal metric making φ_t 1-Lipschitz.

Let a, b be arbitrary points of T_0 . Given $t \geq 0$, a t -subdivision σ between a and b is a sequence $(a = y_0, x_1, y_1, \dots, x_n, y_n, x_{n+1} = b)$ such that $y_i \sim_t x_{i+1}$. When there is no risk of confusion, we will simply say *subdivision*. A subdivision is *straight* if all points belong to $[a, b]$ and lie in the indicated order.

The *flesh* of σ consists of the segments $I_i = [x_i, y_i]$. The *jumps* of σ are the segments $J_i = [y_i, x_{i+1}]$. The *length* $|\sigma|$ of σ is the total length of its flesh: $|\sigma| = \sum_{i=1}^n |I_i|$.

Note that the concatenation of the paths $f(I_i)$ gives a path joining $f(a)$ to $f(b)$ in T_∞ . This implies the inequality $|\sigma| \geq d_\infty(f(a), f(b))$. Similarly, the concatenation of the paths $\varphi_t(I_i)$ joins $\varphi_t(a)$ to $\varphi_t(b)$, so to make φ_t 1-Lipschitz we need to choose a metric on T_t so that $d(\varphi_t(a), \varphi_t(b)) \leq |\sigma|$. The metric d_t which we take on T_t will be defined as $d_t(\varphi_t(a), \varphi_t(b)) = \inf_\sigma |\sigma|$, where the infimum is taken over all t -subdivisions between a and b . We first show that this infimum is achieved.

Lemma 3.1 (Admissible subdivisions). *Given $a, b \in T_0$ and $t \in \mathbb{R}^+$, there exists a straight t -subdivision σ_0 between a and b such that $|\sigma_0| \leq |\sigma|$ for every t -subdivision σ between a and b .*

Such a subdivision will be called *admissible* (or t -admissible if there is a risk of confusion).

Proof. All subdivisions considered in this proof are t -subdivisions between a and b . Define the *complexity* of a subdivision as $c(\sigma) = (|\sigma|, \sum_{i=0}^n |J_i|)$, with $\sum |J_i|$ the total length of the jumps. Complexities are compared in lexicographic order.

Let K be a finite subtree of T_0 containing a and b . We say that a subdivision is contained in K if all points in the sequence are in K . We first show that there exists a number M_K such that, given $\sigma \subset K$, there exists $\sigma' \subset K$ consisting of at most M_K points, with $c(\sigma') \leq c(\sigma)$.

Since f is a morphism, we may cut K into finitely many arcs, all of which are mapped injectively into T_∞ . If two non-consecutive points of σ belong to the same arc, we may decrease n by removing points in-between. This does not increase complexity, because $|\mu| \geq d_\infty(f(x), f(y))$ for any subdivision μ between x and y , and shows the existence of M_K .

Since \sim_t is a closed equivalence relation, a simple compactness argument on the set of subdivisions of cardinality at most M_K shows that, among all subdivisions contained in K , we can find σ_K with minimal complexity.

We now prove that σ_K is straight. Suppose it is not. Exchanging a and b if needed, we may assume that there is an overlap between $[x_i, y_i]$ and $[y_i, x_{i+1}]$. Since x_{i+1} and y_i have the same image p by f , and f is a morphism, the set $(y_i, x_{i+1}] \cap f^{-1}(p)$ is nonempty. Let z be its point closest to y_i . We have $z \sim_t x_{i+1}$, because $\tau(z, x_{i+1}) \leq \tau(y_i, x_{i+1}) \leq t$.

Note that the image of (y_i, z) by f is contained in one component of $T_\infty \setminus \{p\}$. Denoting by $y_{i,\varepsilon}$ (resp. z_ε) the point of (y_i, z) at distance ε from y_i (resp. z), this implies $y_{i,\varepsilon} \sim_t z_\varepsilon$ for $\varepsilon > 0$ small enough. For such an ε , consider σ_ε obtained from σ_K by changing $(\dots x_i, y_i, x_{i+1} \dots)$ into $(\dots x_i, y_{i,\varepsilon}, z_\varepsilon, z, x_{i+1} \dots)$. We have $|\sigma_\varepsilon| = |\sigma_K|$, but $c(\sigma_\varepsilon) < c(\sigma_K)$, a contradiction. This shows that σ_K is straight.

It is now easy to conclude. Define σ_0 minimizing complexity among all subdivisions contained in $[a, b]$. Given any subdivision σ , choose K containing σ . Then $|\sigma| \geq |\sigma_K|$, and $|\sigma_K| \geq |\sigma_0|$ because σ_K is straight. \square

We define $\delta_t(a, b)$ as $|\sigma|$, with σ a t -admissible subdivision between a and b . Clearly, $\delta_t(a, b) = 0$ if and only if $a \sim_t b$. Moreover, δ_t is a pseudo-metric on T_0 . It follows that δ_t defines a metric d_t on T_t , with $d_t(\varphi_t(a), \varphi_t(b)) = \delta_t(a, b)$.

Note that d_0 is just the original distance on T_0 . Given a and b , the function $t \mapsto \delta_t(a, b)$ is non-increasing and equals $\delta_\infty(a, b) = d_\infty(f(a), f(b))$ for $t \geq \frac{d_0(a, b)}{2}$.

Consider an admissible subdivision σ between a and b . Its flesh arcs map isometrically into T_t , and the concatenation of their images is a path of d_t -length $|\sigma| = \delta_t(a, b)$ joining $\varphi_t(a)$ to $\varphi_t(b)$. This path is therefore geodesic. The fact that T_t is an \mathbb{R} -tree will imply that this is *the* arc joining $\varphi_t(a)$ to $\varphi_t(b)$.

Lemma 3.2. *The metric space T_t is an \mathbb{R} -tree.*

Proof. Since T_t is connected, it suffices to show that the distances between any four points satisfy the 0-hyperbolicity inequality [1, Theorem 3.17]. If not, we can find $\theta > 0$, and $a_0, b_0, c_0, d_0 \in T_0$ satisfying the following equation I_θ :

$$\delta_t(a_0, b_0) + \delta_t(c_0, d_0) = \theta + \max(\delta_t(a_0, c_0) + \delta_t(b_0, d_0), \delta_t(a_0, d_0) + \delta_t(b_0, c_0)).$$

Fix such $\theta, a_0, b_0, c_0, d_0$. Among all quadruples $Q = (a, b, c, d)$ contained in the convex hull of $\{a_0, b_0, c_0, d_0\}$ and satisfying I_θ , choose one for which the total length of the convex hull $C(Q)$ is minimal. We may assume that a is a terminal point of $C(Q)$. Consider three admissible subdivisions (a, x_1, y_1, \dots) between a and the other three points.

Each of these subdivisions has positive length, because I_θ forces a, b, c, d to have distinct images in T_t . We may therefore assume that the first flesh interval $I_1 = [x_1, y_1]$ is non-degenerate in all three subdivisions. Then $x_1 \neq a$ in at least one subdivision, since otherwise we could decrease the length of $C(Q)$ by moving a (without losing the equation I_θ). Thus $a \sim_t a'$ for some $a' \neq a$ in $C(Q)$. Since a', b, c, d satisfies I_θ , one can decrease the total length of $C(Q)$, a contradiction. \square

Lemma 3.3. *The maps $\varphi_t : T_0 \rightarrow T_t$ and $\psi_t : T_t \rightarrow T_\infty$ are morphisms.*

Proof. They are obviously surjective, since f is surjective (being a morphism). Now let I be any arc on which f is isometric. Since φ_t and ψ_t are both 1-Lipschitz, φ_t is isometric in restriction to I , and ψ_t is isometric in restriction to $\varphi_t(I)$. This implies that φ_t is a morphism. Furthermore ψ_t is also a morphism, since any arc of T_t is contained in a finite union of images $\varphi_t(I)$. \square

Assume that T_0 and T_∞ are endowed with isometric actions of a group G , and f is equivariant. It is clear that T_t inherits an isometric action of G and that φ_t and ψ_t are equivariant. We will see that T_t is simplicial if T_0 and T_∞ are, but minimality of T_0 and T_∞ as G -trees does not imply minimality of T_t (see 3.3).

3.2. Continuity.

Fix a discrete group G . Let \mathcal{A} be the space of G -trees, i.e. \mathbb{R} -trees T with an isometric action of G . Distance will always be denoted by d . Two trees are

considered equal if there is an equivariant isometry between them. The set \mathcal{A} is equipped with the *equivariant Gromov-Hausdorff topology*.

Recall that a fundamental system of neighborhoods for $T \in \mathcal{A}$ is given by the sets $V_T(X, A, \varepsilon)$, with $X \subset T$ and $A \subset G$ finite sets, and $\varepsilon > 0$. By definition, T' is in $V_T(X, A, \varepsilon)$ if and only if there exists a “lifting” map $x \mapsto x'$ from X to T' such that $|d(x, gy) - d(x', gy')| < \varepsilon$ for every $x, y \in X$ and $g \in A$.

Now consider the space \mathcal{M} of (surjective) G -equivariant morphisms $f : T \rightarrow T_\infty$ between G -trees. Its topology is defined by neighborhoods $W_f(X, A, \varepsilon)$, with $f' : T' \rightarrow T'_\infty$ in $W_f(X, A, \varepsilon)$ if and only if there exists $x \mapsto x'$ as above, with the extra requirement $|d(f(x), f(gy)) - d(f'(x'), f'(gy'))| < \varepsilon$. Note that we define the same topology if we drop g from the last requirement, and that the source and target maps from \mathcal{M} to \mathcal{A} are continuous.

The construction given in the previous section associates morphisms $\varphi_t = \Phi(f, t)$ and $\psi_t = \Psi(f, t)$ to any $f \in \mathcal{M}$ and $t \in [0, \infty]$.

Proposition 3.4. *The maps $\Phi, \Psi : \mathcal{M} \times [0, \infty] \rightarrow \mathcal{M}$ are continuous.*

The main step in the proof is the following lemma:

Lemma 3.5. *Fix $f \in \mathcal{M}$, $\varepsilon > 0$, $t \in [0, \infty)$, and $a, b \in T$. There exist $\alpha > 0$ and a finite set $F \subset [a, b]$ containing a and b , such that, if $f' \in W_f(F, \{1\}, \alpha)$ and $|s - t| < \alpha$, then $|\delta_s(a', b') - \delta_t(a, b)| < \varepsilon$.*

In this statement, a' and b' are lifts of a, b to the source of f' provided by the definition of $W_f(F, \{1\}, \alpha)$, and δ_s is computed with respect to f' .

The lemma immediately implies the continuity of Φ on $\mathcal{M} \times [0, \infty)$. Continuity for $t = \infty$ is clear because, if $f' \in W_f(\{a, b\}, \{1\}, \alpha)$, then $\delta_s(a', b')$ is constant for $t \geq \frac{1}{2}(d(a, b) + \alpha)$. Continuity of Ψ is also clear once we know that of Φ , so 3.4 follows from 3.5. There remains to prove Lemma 3.5.

Proof of Lemma 3.5. In the following proof, C, C_1, C_2 will denote universal constants (which could be made explicit).

Since $\delta_t(a, b)$ is determined by the restriction of f to $[a, b]$, we may assume $T = [a, b]$. We may also forget about G , so we simply write $W_f(F, \alpha)$. Let E consist of a, b , and all points where f folds. Consider $f' \in W_f(E, \alpha)$.

Let $x \mapsto \bar{x}$ be the linear map taking $[a, b]$ to $[a', b']$ (so $\bar{a} = a'$ and $\bar{b} = b'$). It distorts distances by at most α (by which we mean $|d(x, y) - d(\bar{x}, \bar{y})| \leq \alpha$). We claim that

$$|d(f'(\bar{x}), f'(\bar{y})) - d(f(x), f(y))| \leq C\alpha \quad (1)$$

for all $x, y \in [a, b]$. Such an inequality clearly holds for points of E because \bar{x} is $(C_1\alpha)$ -close to x' for $x \in E$. It is also true for arbitrary points because $x \mapsto f'(\bar{x})$ distorts distances by at most $(C_2\alpha)$ on any interval on which f does not fold. We shall assume $C \geq 1$.

We first prove the lemma when $a \sim_t b$ and no point of (a, b) has the same image in T_∞ as a and b . In this case we take $F = E$. Choose θ such that the θ -ball around $f(a)$ contains no vertex of the finite tree $f([a, b])$ different from $f(a)$,

and no image of a point where f folds. Choose α with $10C\alpha < \varepsilon$ and $4C\alpha < \theta$. Consider $f' \in W_f(F, \alpha)$ and $s \geq t - \alpha$.

Viewing $[a, b]$ as a subinterval of \mathbb{R} , we have $f(a + \beta) = f(b - \beta)$ for $0 \leq \beta \leq \theta$. The intervals $I = [a + 2C\alpha, a + 4C\alpha]$ and $J = [b - 4C\alpha, b - 2C\alpha]$ are mapped isometrically onto the same arc in T_∞ . Using (1), we can then find $a_1 \in \overset{\circ}{I}$ and $b_1 \in \overset{\circ}{J}$ with $f'(\bar{a}_1) = f'(\bar{b}_1)$. Note that $f([a_1, b_1])$ is disjoint from $f([a, a + 2C\alpha])$. In particular, since $\tau(a, b) \leq t$, every point of $f([a_1, b_1])$ is $(t - 2C\alpha)$ -close to $f(a_1)$. By (1), we have $\tau(\bar{a}_1, \bar{b}_1) \leq t - 2C\alpha + C\alpha \leq s$ and therefore $\bar{a}_1 \sim_s \bar{b}_1$. We deduce

$$\delta_s(\bar{a}, \bar{b}) \leq \delta_s(\bar{a}, \bar{a}_1) + \delta_s(\bar{a}_1, \bar{b}_1) + \delta_s(\bar{b}_1, \bar{b}) \leq 10C\alpha < \varepsilon$$

since $\delta_s(\bar{a}, \bar{a}_1) \leq d(\bar{a}, \bar{a}_1) \leq d(a, a_1) + \alpha \leq 4C\alpha + \alpha \leq 5C\alpha$. The lemma is proved in this case.

When $a \sim_t b$ but $A = f^{-1}(f(a))$ is larger than $\{a, b\}$, we take $F = E \cup A$ and we prove the lemma by applying the previous argument on each subinterval of $[a, b]$ bounded by points of A .

In the general case, we fix a t -admissible subdivision $\sigma = (a, x_1, y_1, \dots, x_n, y_n, b)$ between a and b . Define F as the union of E and all preimages $f^{-1}(f(x))$, for x a point of σ . Let $f' \in W_f(F, \alpha)$ and $s \in (t - \alpha, t + \alpha)$. We will force $|\delta_s(\bar{a}, \bar{b}) - \delta_t(a, b)| < \varepsilon$ by choosing α small enough.

We have $\delta_t(a, b) = \sum_{i=1}^n d(x_i, y_i)$, and

$$\delta_s(\bar{a}, \bar{b}) \leq \sum_{i=1}^n d(\bar{x}_i, \bar{y}_i) + \sum_{i=0}^n \delta_s(\bar{y}_i, \bar{x}_{i+1}).$$

Furthermore, $d(\bar{x}_i, \bar{y}_i) \leq d(x_i, y_i) + \alpha$. Since $y_i \sim_t x_{i+1}$, we have seen that we can choose α small enough so that $\delta_s(\bar{y}_i, \bar{x}_{i+1}) \leq \varepsilon/2(n+1)$. If furthermore $n\alpha < \varepsilon/2$ we obtain $\delta_s(\bar{a}, \bar{b}) \leq \delta_t(a, b) + \varepsilon$.

We now get a lower bound for $\delta_s(\bar{a}, \bar{b})$. We may assume that f does not fold on $I_i = [x_i, y_i]$ (if $t = 0$, we may have to refine σ). The segment $[(\bar{x}_i)_s, (\bar{y}_i)_s] \subset T'_s$ then has length at least $|I_i| - C\alpha$. As above, we choose α so that $(\bar{y}_i)_s$ and $(\bar{x}_{i+1})_s$ are $(\varepsilon/2(n+1))$ -close.

We shall prove that the intersection between $[(\bar{x}_i)_s, (\bar{y}_i)_s]$ and $[(\bar{x}_{i+1})_s, (\bar{y}_{i+1})_s]$ has length at most $3C\alpha$. For ε and $C\alpha$ small with respect to $\min_i |I_i|$ we can then write

$$\delta_s(\bar{a}, \bar{b}) \geq \sum_{i=1}^n (|I_i| - C\alpha) - \sum_{i=0}^n \left(\frac{\varepsilon}{2(n+1)} + 3C\alpha \right) \geq \delta_t(a, b) - 4(n+1)C\alpha - \frac{\varepsilon}{2}.$$

This gives the required inequality provided $\alpha < \frac{\varepsilon}{8(n+1)C}$.

If the intersection is bigger than $3C\alpha$, we can find $\bar{u} \in [(\bar{x}_i)_s, (\bar{y}_i)_s]$ and $\bar{v} \in [(\bar{x}_{i+1})_s, (\bar{y}_{i+1})_s]$ with $\bar{u}_s = \bar{v}_s$, and both $d(\bar{u}, \bar{y}_i)$, $d(\bar{v}, \bar{x}_{i+1})$ bigger than $3C\alpha$. We shall reach a

contradiction by considering the corresponding points $u \in I_i$ and $v \in I_{i+1}$. Let $p = f(y_i) = f(x_{i+1}) \in T_\infty$.

We have $d(f(u), p) = d(u, y_i) \geq 3C\alpha - \alpha \geq 2C\alpha$. Similarly, $d(f(v), p) \geq 2C\alpha$. On the other hand, $d(f(u), f(v)) \leq C\alpha$. This implies that the arcs $f(I_i)$ and $f(I_{i+1})$ belong to (the closure of) the same component \mathcal{C} of $T_\infty \setminus \{p\}$. In particular, the points $y_i - \kappa$ and $x_{i+1} + \kappa$ have the same image in T_∞ for $\kappa > 0$ small. Since they don't have the same image in T_t , there exists $z \in [y_i, x_{i+1}]$ such that $f(z)$ is at distance exactly t from p , and in a component of $T_\infty \setminus \{p\}$ other than \mathcal{C} . This implies that the distance from $f(z)$ to $f(u)$ is at least $t + 2C\alpha$, and therefore by (1) the point $f'(\bar{z})$ has distance at least $t + C\alpha$ from $f'(\bar{u}) = f'(\bar{v})$. But $\bar{u}_s = \bar{v}_s$, so $s \geq t + C\alpha \geq t + \alpha$, a contradiction.

This completes the proof of Lemma 3.5, hence also of Proposition 3.4. \square

3.3. Additional properties of the semi-flow.

The material from this section will not be used in an essential way in the rest of the paper, as alternative arguments will be provided (see Remark 4.1).

Let $f : T_0 \rightarrow T_\infty$ be a fixed (surjective) morphism between G -trees, with G finitely generated.

Simplicial trees.

Proposition 3.6. *If $f : T_0 \rightarrow T_\infty$ is a morphism between metric simplicial G -trees with finitely many orbits of edges, then the tree T_t is simplicial for every $t > 0$.*

Proof. After subdividing, we may assume that each edge of T_0 is mapped isometrically onto an edge of T_∞ . Fix t . Let $S_\infty \subset T_\infty$ consist of all vertices, and all points lying at distance exactly t from some vertex. Let S_0 and S_t denote the preimages of S_∞ in T_0 and T_t . Since there are finitely many orbits of edges in T_∞ , all three sets S_0, S_t, S_∞ are closed and meet a given finite subtree in finitely many points.

We now show that any point $u \in T_t$ not in S_t is regular: its complement in T_t has only two components. This implies that T_t is simplicial, with vertex set contained in S_t .

Consider $x, y \in T_0$ mapping onto u . They belong to open edges of T_0 . The identification time $\tau(x, y)$ is at most t . If $\tau(x, y) = t$, consider p at distance t from $f(x)$ in $f([x, y])$. It is a vertex of T_∞ , since f is an immersion away from vertices. This implies $\psi_t(u) \in S_\infty$, contradicting $u \notin S_t$. We deduce $\tau(x, y) < t$, and there exist open neighborhoods of x and y in T_0 with the same image in T_t . It follows that u is regular. \square

Minimality.

Let T be an \mathbb{R} -tree with an isometric action of G . Recall that T (or the action) is called minimal if there is no proper G -invariant subtree. If the action of G on T is non-trivial (there is no global fixed point), there is a unique minimal subtree $T_{\min} \subset T$; it consists of all axes of hyperbolic elements of G (see [8]).

A *terminal point* of T is a point x which is not contained in any open interval (equivalently, $T \setminus \{x\}$ is connected). If T is minimal, it contains no terminal point (because $T \setminus G.x$ is a proper invariant subtree if x is terminal).

The action of G on T is *finitely supported* if there exists a finite subtree K such that every segment of T is covered by finitely many translates of K . A simplicial action is finitely supported if and only if there are finitely many orbits of edges. A minimal action of a finitely generated group on an \mathbb{R} -tree is finitely supported.

Lemma 3.7. *A non-trivial finitely supported G -tree T is minimal if and only if it has no terminal point.*

Proof. We have seen that there is no terminal point in a minimal tree. We now assume that T is not minimal and we find a terminal point. Let T_{\min} be the minimal subtree. Consider a finite tree K such that $G.K = T$. Then $K \cap T_{\min}$ is a convex subset of K , distinct from K , so there is at least one endpoint of K which is not contained in T_{\min} . Consider an endpoint x of K whose distance d to T_{\min} is maximal. Clearly, any point of T is at distance at most d from T_{\min} . This implies that x is terminal in T . \square

Let $f : T_0 \rightarrow T_\infty$ be as above. Minimality of T_0 and T_∞ as G -trees does not imply minimality of T_t . For instance, if T_0, T_∞ and f are simplicial, and f maps all edges incident to some vertex $v \in T_0$ onto the same edge of T_∞ , then $\varphi_t(v)$ is a terminal vertex of T_t for t small enough. This is essentially the only way in which T_t may fail to be minimal.

Definition 3.8. *A morphism of \mathbb{R} -trees $f : T_0 \rightarrow T_\infty$ satisfies the minimality condition if, for all $x \in T_0$, there exists an open interval containing x on which f is one-to-one. In particular, T_0 and T_∞ have no terminal point, and are therefore minimal if they are finitely supported.*


Remark 3.9. One can easily prove that, if $f : T_0 \rightarrow T_\infty$ is a morphism between minimal simplicial G -trees, then one can construct another morphism $g : T_0 \rightarrow T_\infty$ satisfying the minimality condition.

The minimality condition clearly implies minimality of T_t :

Proposition 3.10. *Let $f : T_0 \rightarrow T_\infty$ be a morphism between G -trees. Assume that T_0 is finitely supported and f satisfies the minimality condition. Then for all $t \geq 0$ the tree T_t is finitely supported, and the morphisms φ_t, ψ_t satisfy the minimality condition. In particular, T_t is minimal. \square*

The semi-flow property.

Proposition 3.11. *For all $s, t \in \mathbb{R}$, we have $\Psi((\Psi(f, t), s)) = \Psi(f, s + t)$ and $\Phi(\Psi(f, t), s) \circ \Phi(f, t) = \Phi(f, s + t)$: folding T_t along distance s gives the same result as folding T_0 along $s + t$.*



Proof. Fix $s, t \geq 0$, and denote by a_t the image of a point $a \in T_0$ in $T_t = T_t(f)$. We have to show $\delta_{s+t}(a, b) = \delta_s(a_t, b_t)$ (with δ_s computed with respect to $\Psi(f, t) : T_t \rightarrow T_\infty$). We are going to prove the equality

$$\tau(x_t, y_t) = \max(0, \tau(x, y) - t), \quad (2)$$

valid whenever $f(x) = f(y)$ (with $\tau(x_t, y_t)$ also computed with respect to $\Psi(f, t)$).

We first check that (2) implies the lemma. By (2), the projection onto T_t of an $(s+t)$ -subdivision σ between a and b is an s -subdivision σ_t between a_t and b_t . Since $|\sigma_t| \leq |\sigma|$, one gets $\delta_{s+t}(a, b) \geq \delta_s(a_t, b_t)$. Conversely, consider an s -subdivision σ_t between a_t and b_t in T_t , and lift it arbitrarily to σ between a and b . By (2), σ is an $(s+t)$ -subdivision. If $[x_i, y_i]$ are the arcs in the flesh of σ , then $\delta_{s+t}(a, b) \leq \sum_i \delta_{s+t}(x_i, y_i) \leq \sum_i \delta_t(x_i, y_i) = |\sigma_t|$. This implies $\delta_{s+t}(a, b) \leq \delta_s(a_t, b_t)$.

Let us now prove (2). Fix a t -admissible subdivision σ_0 between x and y . The set $A = \Psi(f, t)([x_t, y_t]) \subset T_\infty$ is the image by f of the flesh of σ_0 . To get the image of $[x, y]$ by f , one has to add the image of the jumps $[y_i, x_{i+1}]$. These are loops based at points of A , each contained in the t -ball around its basepoint. This shows $\tau(x, y) \leq \tau(x_t, y_t) + t$, because A is contained in the $\tau(x_t, y_t)$ -ball around $f(x)$.

We conclude by proving the opposite inequality $\tau(x_t, y_t) \leq \tau(x, y) - t$, when $\tau(x, y) \geq t$ (note that $x_t = y_t$ when $\tau(x, y) \leq t$). Let $B_\infty \subset T_\infty$ be the set of points $p \in f([x, y])$ such that the segment $[f(x), p]$ may be extended geodesically by a segment $[p, q] \subset f([x, y])$ of length t . Note that B_∞ is contained in the $(\tau(x, y) - t)$ -ball around $f(x)$, so we need only prove $\Psi(f, t)([x_t, y_t]) \subset B_\infty$.

So consider B_t , the preimage of B_∞ in $[x_t, y_t]$, and assume that $B_t \subsetneq [x_t, y_t]$. The set B_t is closed, and contains $\{x_t, y_t\}$ because $\tau(x, y) \geq t$. Consider a component (u_t, v_t) of $[x_t, y_t] \setminus B_t$, and choose preimages u_0, v_0 of u_t, v_t in $[x, y] \subset T_0$, with $[u_0, v_0]$ minimal (for inclusion) among all possible choices. Note that $f(u_0) = f(v_0)$ and $\Phi(f, t)((u_0, v_0)) \subset (u_t, v_t)$. Since $f((u_0, v_0)) = \Psi(f, t)((u_t, v_t))$ does not meet the component of $T_\infty \setminus \{f(u_0)\}$ that contains $f(x)$, it is contained in the t -ball around $f(u_0)$. We get $u_0 \sim_t v_0$ and $u_t = v_t$, a contradiction. \square

4. THE OUTER SPACE

We fix a finitely generated group G , decomposed as $G = G_1 * \cdots * G_p * F_k$ with each G_i non-trivial. We require $p \geq 1$ and $p + k \geq 2$, but in this section G_i may be \mathbb{Z} or a nontrivial free product. In particular, the space constructed in [23] to study symmetric automorphisms is a special case of the space PO defined below.

All trees considered here will be minimal simplicial G -trees, up to equivariant isometry. All maps between trees will be G -equivariant (hence onto). Unless otherwise indicated, we assume that there is no *redundant vertex*: if v has degree 2, then v is the fixed point of an element of G exchanging the two edges incident to v .

Let \mathcal{O} be the space of *metric simplicial G -trees* T such that:

- The action of G on T is minimal, with trivial edge stabilizers.
- For each i , there is exactly one orbit of vertices with stabilizer conjugate to G_i .
- All other points have trivial stabilizer.

Via Bass-Serre theory, an element of \mathcal{O} may also be viewed as a *marked metric graph of groups*. It is a finite graph of groups Γ , with an isomorphism from $\pi_1(\Gamma)$ to G (the marking), well-defined up to composition with inner automorphisms. Edge groups are trivial. There is one vertex v_i with group conjugate to G_i for each i , and all other vertex groups are trivial. Edges are assigned a positive length. Minimality of the action on T translates to the fact that every terminal vertex is a v_i .

The quotient of \mathcal{O} by the natural action of $(0, \infty)$ (defined by rescaling, i.e. multiplying all lengths by the same number) is the *projectivized space* $P\mathcal{O}$ (the outer space of G). We equip \mathcal{O} with the equivariant Gromov-Hausdorff topology, and $P\mathcal{O}$ with the quotient topology.

The action of G on T defines a length function $\ell_T : G \rightarrow \mathbb{R}$, by $\ell_T(g) = \min_{x \in T} d(x, gx)$. As actions in \mathcal{O} are minimal with non-abelian length function, assigning to $T \in \mathcal{O}$ its length function defines homeomorphisms from \mathcal{O} onto its image in \mathbb{R}^G , and from $P\mathcal{O}$ onto its image in $P\mathbb{R}^G$ [24].

We often identify \mathcal{O} and $P\mathcal{O}$ with their images. The topology on \mathcal{O} and $P\mathcal{O}$ is commonly called the length function topology, or the *axes topology*.

We now consider the simplicial structure on \mathcal{O} and $P\mathcal{O}$. The metric on a tree $T \in \mathcal{O}$ is determined by finitely many positive numbers, one length for each orbit of edges (equivalently, one length for each edge of Γ). The set of trees obtained from T by varying these numbers will be called the (open simplicial) *cone* containing T . In other words, a cone consists of equivariantly homeomorphic trees (trees with the same underlying simplicial tree).

The projection of a cone to $P\mathcal{O}$ is an *open simplex*. Its closure in $P\mathbb{R}^G$ is a closed simplex Σ , whose points may be viewed as projectivized length assignments on the set of edges of Γ . The intersection of Σ with $P\mathcal{O}$ will be called a *closed simplex* of $P\mathcal{O}$ (a point of Σ is in $P\mathcal{O}$ if and only if the union of edges of length 0 is a forest, and every component of the forest contains at most one v_i , so a closed simplex of $P\mathcal{O}$ is a finite union of open simplices).

Just like Culler-Vogtmann's outer space, $P\mathcal{O}$ is not a simplicial complex: simplices may have faces "at infinity". If one insists on having a simplicial complex, one replaces $P\mathcal{O}$ by its barycentric spine (there is an equivariant deformation retraction of $P\mathcal{O}$ onto its spine, as in [9, 23]).

The *weak topology* on $P\mathcal{O}$ is defined in the usual way: a set is closed if and

only if its intersection with every closed simplex is closed. This topology does not coincide with the axes topology, because the simplicial structure is not locally finite (see [23]). The two topologies have the same restriction, however, on any finite union of simplices (see [15] for a detailed study).

Our goal now is to show that $P\mathcal{O}$ is contractible, both in the axes topology and in the weak topology. Our main tool will be the semi-flow.

Remark 4.1. \mathcal{O} is invariant under the semi-flow, in the following sense. Let $f : T_0 \rightarrow T_\infty$ be a morphism between trees in \mathcal{O} , and $t > 0$. The tree T_t is simplicial by Proposition 3.6 (this will also follow from Lemma 4.3). It clearly has the correct edge and vertex stabilizers, because both T_0 and T_∞ do.

Furthermore, T_t is minimal if f maps each edge isometrically and all vertex stabilizers of T_0 are non-trivial. This follows from Proposition 3.10, but here is a direct argument. Given an edge vw of T_0 , choose a non-trivial element g_v (resp. g_w) fixing v (resp. w). The translation axis of $g_v g_w$ in T_t contains the images of v and w , so the minimal subtree of T_t contains every vertex and T_t is minimal.

Theorem 4.2. *The space $P\mathcal{O}$ is contractible in the axes topology.*

Proof. We use the semi-flow to construct a contraction in \mathcal{O} . We first define a basepoint (or rather a basecone) in \mathcal{O} (see figure 2).




FIGURE 2. The base point of outer space

Choose a free basis t_1, \dots, t_k of F_k , and let $T_0 \in \mathcal{O}$ be described by Figure 2. Formally, T_0 is such that every vertex has non-trivial stabilizer, and the vertex stabilized by G_1 is adjacent to the vertices stabilized by $G_2, \dots, G_p, t_1 G_1 t_1^{-1}, \dots, t_p G_1 t_p^{-1}$. This uniquely defines T_0 as a simplicial non-metric G -tree, as we haven't specified edge-lengths yet.

Given $T \in \mathcal{O}$, we define an equivariant map $f : T_0 \rightarrow T$ by mapping every vertex of T_0 to the unique vertex of T with the same stabilizer, and extending linearly on every edge.

There is a unique edge-length assignment on T_0 such that f is isometric on every edge (in particular, f is a morphism). We thus associate to $T \in \mathcal{O}$ a metric tree $T_0(T) \in \mathcal{O}$ and a morphism $f_T : T_0(T) \rightarrow T$. As a simplicial tree, $T_0(T)$ does

not depend on T and thus lies in the cone consisting of all metric simplicial trees obtained by varying lengths on T_0 .

Let $T_t(T)$ be the tree defined by applying the semi-flow to f_T . It belongs to \mathcal{O} (see Remark 4.1).

We can now consider the map $\rho : (T, t) \mapsto T_t(T)$, from $\mathcal{O} \times [0, \infty]$ to \mathcal{O} . Once we know that ρ is continuous, we will get a continuous family of paths connecting T to $T_0(T)$, thus providing a deformation retraction of \mathcal{O} onto a cone. This will prove contractibility of \mathcal{O} .

Let us check that ρ is continuous (in the axes topology). The edge-length assignment on T_0 depends continuously on T : if $g_i \in G_i$ is non-trivial, the edge of Γ between v_1 and v_i has length $\frac{1}{2}\ell_T(g_1g_i)$, and the loop associated to t_j has length $\frac{1}{2}\ell_T(g_1t_jg_1t_j^{-1})$. More generally, if x, y are vertices of T_0 and g_x, g_y are non-trivial elements fixing them, then $d(f_T(x), f_T(y)) = \frac{1}{2}\ell_T(g_xg_y)$. It follows that $T_0(T)$ and f_T depend continuously on T , so ρ is continuous because the semi-flow is continuous (Proposition 3.4).

To prove contractibility of $P\mathcal{O}$, we simply reparametrize ρ so that it descends to $\hat{\rho} : P\mathcal{O} \times [0, \infty] \rightarrow P\mathcal{O}$. For instance, we may take $\tilde{\rho}(T, t) = \rho(T, \ell_T(h)t)$, with $h \in G$ chosen so that $\ell_T(h) > 0$ for every $T \in \mathcal{O}$ (take $h = g_1g_2$ if $p \geq 2$ and $h = g_1t_jg_1t_j^{-1}$ if $p = 1$, with $g_i \in G_i$ non-trivial). \square

To prove contractibility in the weak topology, we simply show that $\hat{\rho}$ is also continuous in the weak topology. This requires the following ‘‘finiteness lemma’’.

Lemma 4.3. *Consider two non-metric simplicial G -trees T_0, T , with finitely many orbits of edges. Let $f : T_0 \rightarrow T$ be an equivariant simplicial map (sending every edge onto an edge) such that the preimage $f^{-1}(e)$ of each edge $e \subset T$ is a finite set of edges.*

If f factors through equivariant continuous maps

$$\begin{array}{ccc} T_0 & \xrightarrow{f} & T \\ & \searrow f_1 & \nearrow f_2 \\ & & T' \end{array}$$

with T' an \mathbb{R} -tree and f_1 surjective, then T' is a simplicial tree, and there are only finitely many possibilities for T' up to equivariant homeomorphism (i.e. as a non-metric simplicial G -tree).

Remark. In this lemma, T_0 and T are allowed to have redundant vertices.

Proof. For each open edge e of T , let $K_{e,i}$ ($i \in I_e$) be the (finite) set of connected components of the closure of $f_2^{-1}(e)$. Each $K_{e,i}$ is a finite subtree of T' . If $(e, i) \neq (e', j)$, then either $K_{e,i} \cap K_{e',j} = \emptyset$, or $K_{e,i} \cap K_{e',j}$ consists of one point, which is the image of an endpoint of an edge in $f^{-1}(e)$. In particular, for a given $K_{e,i}$, the set of points occurring as $K_{e,i} \cap K_{e',j}$ is finite.

It follows that T' is simplicial. Furthermore, if a vertex v of T' is not the image of a vertex of T_0 , then there exist edges e_0, e'_0 of T_0 , with the same image in T ,

such that v is an endpoint of $f_1(e_0) \cap f_1(e'_0)$. This implies that there is a uniform bound (depending only on f) for the number of G -orbits of vertices of T' .

Now subdivide the trees so as to make f_1 and f_2 simplicial: we first subdivide T by adding images of vertices of T' , then we subdivide T_0 and T' by adding all preimages of vertices of T . This may create redundant vertices, but the number of orbits of redundant vertices in T and T_0 is uniformly bounded; in particular, there are only finitely many possibilities for the subdivisions of T and T_0 .

We fix these subdivisions, and we show that there are only finitely many possibilities for T' . Let e_1, \dots, e_n be a set of representatives for G -orbits of oriented edges of T (subdivided). In the set E of oriented edges of (the subdivided) T_0 , define $e \sim e'$ if e, e' have the same image in T' . There are only finitely many possibilities for \sim , since it is determined by its restriction to the finite set $\cup_i f^{-1}(e_i)$.

We conclude the proof by showing that \sim completely determines T' (as a non-metric G -tree). Consider the connected graph T_0/\sim obtained from T_0 by identifying e, e' whenever $e \sim e'$. Since T' is obtained from T_0/\sim by identifying vertices, but no edges, simple connectedness of T' implies that $T' = T_0/\sim$. \square

Corollary 4.4. *The space $P\mathcal{O}$ is contractible in the weak topology.*

Proof. We first deduce from Lemma 4.3 that, if S is an open simplex of $P\mathcal{O}$, then the restriction of $\hat{\rho}$ to $S \times [0, \infty]$ meets only finitely many simplices. Let \tilde{S} be the preimage of S in O . Given $T \in \tilde{S}$, consider the map f_T constructed in the proof of Proposition 4.2 and subdivide T_0 so that f_T becomes simplicial (this creates redundant vertices). As a simplicial map (forgetting edge lengths), f_T does not depend on $T \in \tilde{S}$. Two distinct edges of T_0 with the same image in T are in different G -orbits (because T has trivial edge stabilizers), so the finiteness condition in Lemma 4.3 is satisfied. The lemma then implies that for $T \in \tilde{S}$ there are only finitely many possibilities for $T_t(T)$ as a simplicial tree.

We can now prove that $\hat{\rho}$ is continuous in the weak topology. Since a closed simplex Σ of $P\mathcal{O}$ is the union of finitely many open simplices, the restriction of $\hat{\rho}$ to $\Sigma \times [0, \infty]$ is continuous in the weak topology because it is continuous in the axes topology, and the two topologies coincide on any finite union of simplices. By definition of the weak topology, this implies continuity of $\hat{\rho}$. \square

Remark. In [15], we extend these results and prove the contractibility of any *deformation space*. In particular, this includes outer space and spaces of JSJ splittings of finitely presented groups. This extended result also implies that the fixed point set of any finitely generated subgroup F of $\text{Out}(G)$ acting on $P\mathcal{O}$ is contractible or empty (and non-empty if F is finite and solvable).

5. THE ACTION OF $\text{Out}(G)$

We denote by $Z(H)$ the center of a group H . The group of inner automorphisms is $\text{Inn}(H) \simeq H/Z(H)$, and the group of outer automorphisms is $\text{Out}(H) = \text{Aut}(H)/\text{Inn}(H)$. Note that a nontrivial free product has trivial center.

Let $G = G_1 * \cdots * G_p * F_k$ be as above, with $p \geq 1$ and $p + k \geq 2$. We now assume that G_i is freely indecomposable and not isomorphic to \mathbb{Z} . In particular, G is a nontrivial free product and is not free.

Since any automorphism of G maps G_i onto a conjugate of some G_j , there is a natural action of $\text{Out}(G)$ on the contractible complex $P\mathcal{O}$, obtained by precomposing actions on trees by automorphisms. If an element of $P\mathcal{O}$ is viewed as a graph of groups Γ , the action is by changing the marking.

Let S be an open simplex of $P\mathcal{O}$. We will think of it as a simplicial G -tree T (with no metric specified), or as a marked graph of groups Γ . If we forget the marking of Γ , there are only finitely many possibilities for Γ (because of the minimality assumption). This means that there are only finitely many orbits of simplices under the action of $\text{Out}(G)$ on $P\mathcal{O}$.

Let us now describe the stabilizer of S , that is the subgroup $\text{Out}^S(G) \subset \text{Out}(G)$ sending S to itself. It is the group of automorphisms preserving the decomposition of G as a graph of groups given by Γ . Such groups have been studied in [2] and [20].

Let $\text{Out}_0^S(G) \subset \text{Out}^S(G)$ be the finite index subgroup consisting of automorphisms acting trivially on the quotient graph $\Gamma = T/G$. Since the edge groups of Γ are trivial, $\text{Out}_0^S(G)$ has a very simple description (see [20, Prop. 4.2]). Denoting by n_i the degree of the vertex v_i in Γ , it is a direct product $\prod_{i=1}^p M_{n_i}(G_i)$ of groups which fit in exact sequences

$$\{1\} \rightarrow G_i^{n_i}/Z(G_i) \rightarrow M_{n_i}(G_i) \rightarrow \text{Out}(G_i) \rightarrow \{1\} \quad (3)$$

$$\{1\} \rightarrow G_i^{n_i-1} \rightarrow M_{n_i}(G_i) \rightarrow \text{Aut}(G_i) \rightarrow \{1\}, \quad (4)$$

with the center $Z(G_i)$ embedded diagonally into $G_i^{n_i}$ (the groups $M_{n_i}(G_i)$ are denoted by $\text{PMCG}^\partial(G_i)$ in [20], as they are ‘‘pure mapping class groups’’).

The exact sequence (4) is split and $M_{n_i}(G_i)$ is the semi-direct product $G_i^{n_i-1} \rtimes \text{Aut}(G_i)$ associated to the diagonal action of $\text{Aut}(G_i)$ on $G_i^{n_i-1}$. In particular, $M_1(G_i) = \text{Aut}(G_i)$, and $M_2(G_i)$ is the holomorph $\text{Hol}(G_i)$.

Let $\text{Out}'(G) \subset \text{Out}(G)$ be the finite index subgroup consisting of automorphisms mapping each G_i to a conjugate of itself. It maps onto $\text{Out}(G_i)$ in a natural way (since G_i equals its normalizer in G). We also consider the homomorphism $\text{Out}'(G) \rightarrow \text{Out}(F_k)$ obtained by viewing F_k as the quotient of G by the normal subgroup generated by the G_i 's.

Lemma 5.1. *If all groups G_i and $G_i/Z(G_i)$ are torsion-free, then the kernel of $\pi : \text{Out}'(G) \rightarrow \text{Out}(F_k) \times \prod_i \text{Out}(G_i)$ is torsion-free.*

Proof. If not, choose $\Phi \in \ker \pi$ whose order is a prime number p . It fixes a point in $P\mathcal{O}$, since otherwise $\mathbb{Z}/p\mathbb{Z}$ would act freely on $P\mathcal{O}$ and have a finite-dimensional classifying space. This means that Φ belongs to some $\text{Out}^S(G)$.

Consider the action of Φ on the quotient graph of groups Γ corresponding to S . It acts trivially on the topological fundamental group $\pi_1(\Gamma) \simeq F_k$, and sends each vertex v_i to itself. Thus Φ acts trivially on Γ , so $\Phi \in \text{Out}_0^S(G)$.

We have seen that $\text{Out}_0^S(G) = \prod_i M_{n_i}(G_i)$. Since Φ maps trivially into $\text{Out}(G_i)$, it belongs to the product $\prod_i G_i^{n_i}/Z(G_i)$. But $H^n/Z(H)$ is torsion-free if both H and $H/Z(H)$ are (if $h_1^p = \cdots = h_n^p \in Z(H)$ with $p \geq 2$, then $h_i \in Z(H)$ because $H/Z(H)$ is torsion-free, and $h_i = h_j$ because $(h_i h_j^{-1})^p = h_i^p h_j^{-p} = 1$). \square

We can now prove:

Theorem 5.2.

- (i) *Assume that G_i and $G_i/Z(G_i)$ are torsion free, and $\text{Out}(G_i)$ is virtually torsion-free. Then $\text{Out}(G)$ and $\text{Aut}(G)$ are virtually torsion-free.
If, furthermore, G_i , $G_i/Z(G_i)$, and $\text{Out}(G_i)$ have finite virtual cohomological dimension (resp. have a finite index subgroup with a finite classifying space), then so do $\text{Out}(G)$ and $\text{Aut}(G)$.*
- (ii) *If G_i and $\text{Aut}(G_i)$ have finite virtual cohomological dimension (resp. have a finite index subgroup with a finite classifying space), then so do $\text{Out}(G)$ and $\text{Aut}(G)$, provided that they are virtually torsion free.*

Basic facts about virtual cohomological dimension and finite classifying spaces are recalled in section 6 of [23].

Remark. If $p + k \geq 3$, then G_i , $G_i/Z(G_i)$ and $\text{Aut}(G_i)$ are isomorphic to subgroups of $\text{Out}(G)$. In particular, if $\text{Out}(G)$ is virtually torsion free (resp. has finite cohomological dimension), so are G_i , $G_i/Z(G_i)$ and $\text{Aut}(G_i)$.

Proof. Since $\text{Aut}(G)$ is an extension of $\text{Out}(G)$ by $\text{Inn}(G) \simeq G$, it suffices to prove the required results for $\text{Out}(G)$.

The first assertion of (i) follows from Lemma 5.1, since $\text{Out}(F_k)$ is virtually torsion-free.

To prove the second assertion, we consider the action of $\text{Out}(G)$ on the contractible complex $P\mathcal{O}$ (strictly speaking, one should replace $P\mathcal{O}$ by its barycentric spine as in [23]). We have seen that there are only finitely many orbits of simplices, so it suffices to check that all stabilizers $\text{Out}^S(G)$ have the property under consideration. As explained above, $\text{Out}^S(G)$ is virtually a direct product of groups $M_{n_i}(G_i)$. The exact sequences (3) and (4) then imply that $M_{n_i}(G_i)$, hence also $\text{Out}^S(G)$, has the required property.

The proof of (ii) is similar. Note that the groups $M_{n_i}(G_i)$ are virtually torsion-free, because they are subgroups of $\text{Out}(G)$. \square

Here is a slightly stronger result:

Corollary 5.3. *Suppose that each G_i is finitely generated and has a normal subgroup of finite index H_i , with H_i and $H_i/Z(H_i)$ torsion-free, and $\text{Out}(H_i)$ virtually torsion-free. Then $\text{Out}(G)$ is virtually torsion-free. If furthermore H_i and $\text{Aut}(H_i)$ have finite virtual cohomological dimension, then $\text{Out}(G)$ has finite virtual cohomological dimension.*

The proof requires the following fact (see [19], [22]).

Lemma 5.4. *Let Γ be a finitely generated group. Let Δ be a normal subgroup of finite index, with trivial center.*

- (1) *Some finite index subgroup $\text{Aut}_0(\Gamma) \subset \text{Aut}(\Gamma)$ embeds into $\text{Aut}(\Delta)$.*
- (2) *Some finite index subgroup $\text{Out}_0(\Gamma) \subset \text{Out}(\Gamma)$ is isomorphic to the quotient of a subgroup of $\text{Out}(\Delta)$ by a finite subgroup.*

Proof. Let $\text{Aut}_0(\Gamma)$ consist of all automorphisms of Γ mapping Δ to Δ and inducing the identity on Γ/Δ , and let $\text{Out}_0(\Gamma)$ be the image of $\lambda : \text{Aut}_0(\Gamma) \rightarrow \text{Out}(\Gamma)$.

We claim that the natural map from $\text{Aut}_0(\Gamma)$ to $\text{Aut}(\Delta)$ is injective. Assume that α is in the kernel. Given $h \in \Gamma$, define the element $\delta_0(h) \in \Delta$ by $\alpha(h) = h\delta_0(h)$. For all $g \in \Delta$, we have $hgh^{-1} = \alpha(hgh^{-1}) = h\delta_0(h)g\delta_0(h)^{-1}h^{-1}$, so $\delta_0(h) \in Z(\Delta) = \{1\}$ and α is the identity.

The group $\Sigma = \text{Aut}_0(\Gamma)/\text{Inn}(\Delta)$ is therefore isomorphic to a subgroup of $\text{Out}(\Delta)$. The kernel of λ contains $\text{Inn}(\Delta)$ with finite index, so $\text{Out}_0(\Gamma)$ is the quotient of Σ by a finite subgroup. \square

Corollary 5.5. *Let Γ be a finitely generated group. Let Δ be a normal subgroup of finite index, with trivial center. If $\text{Out}(\Delta)$ is virtually torsion-free (resp. has finite virtual cohomological dimension), so is $\text{Out}(\Gamma)$.* \square

Proof of Corollary 5.3. Let Δ the kernel of the map from G to $\prod_i G_i/H_i$. It is a finite free product of groups which are isomorphic to \mathbb{Z} or to an H_i . Corollary 5.3 now follows from Theorem 5.2 and Corollary 5.5. \square

6. APPLICATIONS

6.1. Hyperbolic groups.

Theorem 6.1. *Let G be a hyperbolic group.*

- (1) *If G is torsion-free, then $\text{Out}(G)$ has a finite index subgroup with a finite classifying space.*
- (2) *If G is virtually torsion-free, then $\text{Out}(G)$ has finite virtual cohomological dimension.*

It is not known whether there exist hyperbolic groups which are not virtually torsion-free (see [16]).

Proof. First assume that G is torsion-free. There are two cases (besides $G = \{1\}$ and $G = \mathbb{Z}$). If G is one-ended, then some finite index subgroup of $\text{Out}(G)$ maps onto a direct product of mapping class groups of surfaces, with kernel \mathbb{Z}^n (see [20], [27]). The result is therefore true in this case. If G has infinitely many ends, we simply apply Theorem 5.2. Note that G_i has trivial center, and a finite classifying space (the quotient of the Rips complex).

Assertion (2) follows from Corollary 5.5. \square

6.2. Groups acting freely on \mathbb{R}^n -trees.

In this section, we consider a finitely generated group G admitting a free action on an \mathbb{R}^n -tree. This includes limit groups (also known as finitely generated ω -residually free groups); this is due to Remeslennikov [26], see [14] for a correct statement.

The hypothesis that G acts freely on an \mathbb{R}^n -tree implies that G is torsion-free and commutative transitive (centralizers of nontrivial elements are abelian; see for instance [6]). In particular, if G is not abelian, it has trivial center. Furthermore, G has a finite classifying space, every abelian subgroup is contained in a maximal one, maximal abelian subgroups are finitely generated and malnormal (see [13, 14], and [17, 28] for limit groups).

We first prove:

Theorem 6.2. *If a finitely generated group G acts freely on an \mathbb{R}^n -tree, then $\text{Out}(G)$ has finite virtual cohomological dimension.*

If $n = 1$, then G is a free product of surface groups and free abelian groups by Rips's theorem (see [3], [12]). Since Theorem 6.2 is true for such groups, it is true for G by Theorem 5.2. The proof in the general case will be by induction on n , using the following fact.

Theorem 6.3. *If a finitely generated, freely indecomposable group G acts freely on an \mathbb{R}^n -tree ($n \geq 2$), there exists an $\text{Out}(G)$ -invariant G -tree T with edge stabilizers isomorphic to \mathbb{Z} , and vertex stabilizers acting freely on \mathbb{R}^{n-1} -trees. No edge of T joins two vertices with abelian stabilizers.*

Proof. The tree T is constructed in Proposition 4.2 of [25], as (a small modification of) the cyclic JSJ splitting of G . The only property not proved there is that vertex stabilizers act freely on \mathbb{R}^{n-1} -trees. But a vertex stabilizer G_v is a surface group (and therefore acts freely on an \mathbb{R} -tree) or fixes a point in every G -tree T' with cyclic edge stabilizers [25, Théorème 4.1]. By Theorem 7.1 of [14], there exists such a T' with vertex stabilizers acting freely on \mathbb{R}^{n-1} -trees. It follows that G_v acts freely on an \mathbb{R}^{n-1} -tree. \square

Let V be the vertex set of the graph of groups Γ associated to T . By [20, Proposition 4.2], some finite index subgroup $\text{Out}_1(G) \subset \text{Out}(G)$ fits in an exact sequence

$$1 \rightarrow \mathcal{T} \rightarrow \text{Out}_1(G) \rightarrow \prod_{v \in V} M(G_v) \rightarrow 1,$$

where G_v is the vertex group, $M(G_v)$ is a subgroup of $\text{Out}(G_v)$, and \mathcal{T} is the *group of twists* associated to Γ . The group \mathcal{T} is generated by a finite direct product of centralizers of edge groups in vertex groups (see [20], or the proof below). In particular, \mathcal{T} is a finitely generated abelian group.

The next proposition will show that \mathcal{T} is torsion-free. Assuming this, we complete the proof of Theorem 6.2 as follows.

Proof of Theorem 6.2. First assume that G is freely indecomposable. Apply Theorem 6.3. By the induction hypothesis, all groups $M(G_v)$ have finite virtual cohomological dimension. Since \mathcal{T} is torsion-free with finite cohomological dimension, $\text{Out}_1(G)$ is virtually torsion-free and has finite virtual cohomological dimension. The same is true of $\text{Out}(G)$. If G is a non-trivial free product, we use Theorem 5.2. \square

Proposition 6.4. *Let Γ be a minimal graph of groups decomposition of a commutative transitive group G , with edge groups isomorphic to \mathbb{Z} . Assume that every edge of Γ has at least one endpoint with nonabelian vertex group. Then the group of twists \mathcal{T} is isomorphic to a finite direct product of abelian subgroups of G .*

Proof. We denote by E be the set of oriented edges of Γ , by E_v the set of edges with origin v , by $o(e)$ the origin of e , by $Z_{G_{o(e)}}(G_e)$ the centralizer of the edge group G_e in the vertex group $G_{o(e)}$. We say that a vertex v of Γ is abelian if G_v is abelian. By subdividing, we may assume that each edge joins an abelian vertex to a nonabelian one.

We recall the presentation of \mathcal{T} given in [20, Proposition 3.1]. The group \mathcal{T} is the quotient of $\prod_{e \in E} Z_{G_{o(e)}}(G_e)$ by edge and vertex relations, defined as follows.

For every pair (f, \bar{f}) of opposite edges, we kill the diagonal image of $G_f = G_{\bar{f}}$ in $Z_{G_{o(f)}}(G_f) \times Z_{G_{o(\bar{f})}}(G_{\bar{f}}) \subset \prod Z_{G_{o(e)}}(G_e)$. For every abelian vertex v , we kill the diagonal image of G_v in $\prod_{e \in E_v} Z_{G_v}(G_e) \subset \prod_{e \in E} Z_{G_{o(e)}}(G_e)$ (note that $Z_{G_v}(G_e) = G_v$). There is no relation at non-abelian vertices.

We first reduce to the case where every non-abelian vertex is terminal, simply replacing a non-abelian v with degree $n_v \geq 2$ by n_v terminal vertices each carrying G_v . This may disconnect Γ , but then \mathcal{T} is the direct product of the groups associated to the components of the new graph. We may therefore assume that Γ is a minimal graph of groups of the following form: it consists of a central abelian vertex v , connected to terminal non-abelian vertices v_1, \dots, v_n by edges e_1, \dots, e_n .

First suppose that G_{e_i} is equal to Z_i , its centralizer in G_{v_i} , for all i . Then \mathcal{T} is the quotient of $(G_v)^n$ by G_v (embedded diagonally), so is isomorphic to $(G_v)^{n-1}$. If G_{e_i} is properly contained in Z_i for some i , then by transitive commutativity G_{e_i} maps onto G_v . Furthermore, this may happen only for one value of i , say $i = 1$. The group \mathcal{T} is the quotient of $Z_1 \times \prod_{i>1} G_v$ by the image of G_v , embedded diagonally into the whole product (including the factor Z_1). Since G_{e_1} maps onto G_v , minimality of Γ implies $n \geq 2$, so \mathcal{T} is isomorphic to $Z_1 \times \prod_{i>2} G_v$. \square

Remark. The proposition is true without the assumption on edges of Γ , provided G is not a solvable Baumslag-Solitar group $BS(1, s)$. It also applies to abelian splittings, provided G is not the fundamental group of a graph of groups of the following form: there is only one vertex v , its group G_v is abelian, and all inclusions from edge groups into G_v , except possibly one, are onto.

Our results so far used the cyclic JSJ splitting of G . Using results of [18] and [5] about the abelian JSJ splitting, one gets:

Theorem 6.5. *If G is a limit group, then $\text{Out}(G)$ has a finite index subgroup with a finite classifying space.*

Proof. Since limit groups have finite classifying spaces, Theorem 5.2 lets us assume that G is one-ended. By Theorems 3.13 and 3.17 of [5], it has an $\text{Out}(G)$ -invariant abelian JSJ splitting. The associated graph of groups Γ has three types of vertices: elementary vertices (whose vertex group is a maximal abelian subgroup), surface vertices, rigid vertices. Every edge connects an elementary vertex to a non-elementary vertex.

Since this splitting is $\text{Out}(G)$ -invariant, we may use the exact sequence

$$1 \rightarrow \mathcal{T} \rightarrow \text{Out}_1(G) \rightarrow \prod_{v \in V} M(G_v) \rightarrow 1$$

of [20] as in the proof of Theorem 6.2.

Since not every $\text{Out}(G_e)$ is finite, proving that $\text{Out}_1(G)$ has finite index in $\text{Out}(G)$ now requires the following argument, extending Proposition 2.3 of [20]. An edge e with $\text{Out}(G_e)$ infinite connects an elementary vertex v to a rigid vertex w . Since only finitely many outer automorphisms of G_w extend to automorphisms of G by [18, Theorem 11.1], the group $M(G_w)$ has finite index in the image of ρ_w . Furthermore, every bitwist around e is a twist because the normalizer of G_e in an abelian G_v equals its centralizer.

As before, the group \mathcal{T} is finitely generated and abelian. If v is rigid, then $M(G_v)$ is finite by [18]. If v is a surface vertex, $M(G_v)$ is a surface mapping class group. If G_v is abelian, then $M(G_v)$ is the subgroup of $\text{Aut}(G_v) = GL(n, \mathbb{Z})$ consisting of automorphisms equal to the identity on some finite collection of subgroups; it has a finite index subgroup with a finite classifying space. Since we know that $\text{Out}(G)$ is virtually torsion-free, we conclude that some finite index subgroup has a finite classifying space. \square

Theorem 6.5 is actually valid for a broader class of groups, including groups acting freely on \mathbb{R}^n -trees. This will appear elsewhere, as it requires a general construction of invariant abelian splittings.

REFERENCES

1. R. Alperin, H. Bass, *Length functions of group actions on Λ -trees*, in “Combinatorial group theory and topology (S.M. Gersten, J.R. Stallings, ed.)”, Ann. Math. Studies 111 (1987), Princeton Univ. Press.
2. H. Bass, R. Jiang, *Automorphism groups of tree actions and of graphs of groups*, J. Pure Appl. Algebra **112** (1996), 109–155.
3. M. Bestvina, M. Feighn, *Stable actions of groups on real trees*, Inv. Math. **121** (1995), 287–321.
4. B. Bowditch, *Cut points and canonical splittings of hyperbolic groups*, Acta Math. **180** (1998), 145–186.
5. I. Bumagin, O. Kharlampovich, A. Miasnikov, *Isomorphism problem for finitely generated fully residually free groups*, preprint.
6. I. Chiswell, *Introduction to Λ -trees*, World Scientific Publishing Co., 2001.

7. M. Clay, *Contractibility of deformation spaces of G -trees*, Alg. & Geom. Topology **5** (2005), 1481–1503.
8. M. Culler, J. Morgan, *Group actions on \mathbb{R} -trees*, Proc. London Math. Soc. **55** (1987), 571–604.
9. M. Culler, K. Vogtmann, *Moduli of graphs and automorphisms of free groups*, Invent. Math. **84** (1986), 91–119.
10. M. Forester, *Deformation and rigidity of simplicial group actions on trees*, Geom. & Topol. **6** (2002), 219–267.
11. D.I. Fouxe-Rabinovitch, *On the automorphism group of free products I*, Rec. Math. [Mat. Sbornik] N.S. **8** (1940), 265–276; *II*, Rec. Math. [Mat. Sbornik] N.S. **9** (1941), 183–220.
12. D. Gaboriau, G. Levitt, F. Paulin, *Pseudogroups of isometries of \mathbb{R} and Rips’ theorem on free actions on \mathbb{R} -trees*, Isr. Jour. Math. **87** (1994), 403–428.
13. S. Gross, *Group actions on Λ -trees*, PhD thesis, Hebrew University of Jerusalem (1998).
14. V. Guirardel, *Limit groups and groups acting freely on \mathbb{R}^n -trees*, Geometry and Topology **8** (2004), 1427–1470.
15. V. Guirardel, G. Levitt, *Deformation spaces of trees*, arXiv math.GR/0605545.
16. I. Kapovich, D. Wise, *The equivalence of some residual properties of word-hyperbolic groups*, J. Algebra **223** (2000), 562–583.
17. O. Kharlampovich, A. Myasnikov, *Irreducible affine varieties over a free group. I. Irreducibility of quadratic equations and Nullstellensatz*, J. Algebra **200** (1998), 472–516; *II. Systems in triangular quasi-quadratic form and description of residually free groups*, J. Algebra **200** (1998), 517–570.
18. O. Kharlampovich, A. Myasnikov, *Effective JSJ decompositions*, Contemp. Math. (Groups, languages, algorithms) **378** (2005), 87–212.
19. S. Krstić, *Finitely generated virtually free groups have finitely presented automorphism group*, Proc. London Math. Soc. **64** (1992), 49–69.
20. G. Levitt, *Automorphisms of hyperbolic groups and graphs of groups*, Geom. Dedic. **114** (2005), 49–70.
21. J. Los, M. Lustig, *The set of train track representatives of an irreducible free group automorphism is contractible*, preprint.
22. J. McCool, *The automorphism groups of finite extensions of free groups*, Bull. London Math. Soc. **20** (1988), 131–135.
23. D. McCullough, A. Miller, *Symmetric automorphisms of free products*, Mem. Amer. Math. Soc. **122** (1996).
24. F. Paulin, *The Gromov topology on \mathbb{R} -trees*, Topology and its App. **32** (1989), 197–221.
25. F. Paulin, *Sur la théorie élémentaire des groupes libres*, Séminaire Bourbaki 922, Astérisque **294** (2004), 363–402.
26. V. Remeslennikov, *\exists -free groups as groups with a length function*, Ukrainian Math. J. **44** (1992), 733–738.
27. Z. Sela, *Structure and rigidity in (Gromov) hyperbolic groups and discrete groups in rank 1 Lie groups II*, GAFA **7** (1997), 561–593.
28. Z. Sela, *Diophantine geometry over groups I: Makanin-Razborov diagrams*, Publ. Math. IHES **93** (2001), 31–105.
29. R. Skora, *Deformations of length functions in groups*, Preprint (1989).
30. K. Vogtmann, *Automorphisms of free groups and outer space*, Geom. Dedic. **94** (2002), 1–31.

V.G.: LABORATOIRE ÉMILE PICARD, UMR CNRS 5580, UNIVERSITÉ PAUL SABATIER, 31062 TOULOUSE CEDEX 4, FRANCE.

E-mail address: guirardel@picard.ups-tlse.fr

G.L.: LMNO, UMR CNRS 6139, BP 5186, UNIVERSITÉ DE CAEN, 14032 CAEN CEDEX, FRANCE.

E-mail address: levitt@math.unicaen.fr