R-Trees and Normalization of Pseudogroups

Frank Rimlinger

CONTENTS

Let $\mathcal{G}$ be a pseudogroup defined on a tree Z, and let Γ be a finite set of generators for $\mathcal{G}$. The reduced fundamental group $\bar{\pi}_1(\Gamma)$ of Γ is defined here. I give a new and experimentally inspired proof of a result of Levitt: If $\bar{\pi}_1(\Gamma)$ is a free group, there exists a finite set of generators Ψ for $\mathcal{G}$ such that $\bar{\pi}_1(\Psi)$ is free on the set Ψ. If Ψ has no dead ends, it is an interval exchange.

Like Gaboriau, Levitt and Paulin [Gaboriau et al. 1992], I prove that if G is a finitely presented group acting freely on an $\mathbf{R}$-tree and Γ is a corresponding set of pseudogroup generators, we're in one of the following situations: either G splits as a free product with a noncyclic free abelian summand, or Γ can be reduced to an interval exchange by normalizing and removing a finite number of dead ends, or the process of removing dead ends from Γ does not terminate in a finite number of steps.

INTRODUCTION

A *pseudogroup presentation* Γ defined on a metric space Z is, in general, a collection of local isometries, each of which takes a connected subset of Z onto some other connected subset of Z. Such a set of maps Γ decomposes Z into *orbits*: Two points are in the same orbit if it is possible to pass from one to the other by applying a finite succession of maps of Γ or their inverses. The collection $\mathcal{G}$ of these orbits is the *pseudogroup* generated by Γ. In general, a pseudogroup defined on a space Z is generated by many different presentations, and one wishes to pass from an arbitrary such presentation to a canonical form presentation.

Perhaps the most geometrically satisfying characterization of an $\mathbf{R}$-tree is found at the beginning of [Gaboriau et al. 1992]: An $\mathbf{R}$-tree is a path-connected metric space in which every arc is isometric to an interval of $\mathbf{R}$. From this point of view, a *tree* Z is a connected, simply connected one-complex with a metric that induces the topology of the one-complex structure and gives Z the

This work is based on research partially supported by Fairfield University.

1058-6458/92 $0.50 per page

structure of an **R**-tree. All the fundamental results about **R**-trees needed in this paper are taken from [Rimlinger 1992], which uses results of [Alperin and Bass 1987] developed from [Chiswell 1976], which in turn grew out of the seminal work of Lyndon [1963]. See also [Morgan and Shalen 1984; 1991] and [Shalen 1987] for background information, especially the fact, proved by Morgan and Shalen, that most surface groups act freely on **R**-trees.

Let Γ be a finite set of generating maps for the pseudogroup $\mathcal{G}$—in other words, a finite presentation of $\mathcal{G}$—and let $F(\Gamma)$ be the free group on Γ. Then $F(\Gamma)$ has the structure of a pseudogroup presentation of $\mathcal{G}$. Indeed, a Γ-word is naturally interpreted as a composition of generating maps or their inverses. An element $\alpha \in F(\Gamma)$ is a *cycle* if there exists a point $p \in Z$ such that $\alpha(p) = p$. An element of $F(\Gamma)$ is a *reflection* if it has a fixed point but is not an identity map. (Thus every reflection is a cycle.) We define the *reduced fundamental group* $\bar{\pi}_1(\Gamma)$ to be $F(\Gamma)$ modulo the normal subgroup of $F(\Gamma)$ generated by cycles: compare [Levitt 1990a].

Let Γ be a pseudogroup presentation of $\mathcal{G}$ defined on a tree Z. In Section 2, we show that if $F(\Gamma)$ contains no reflection, $\bar{\pi}_1(\Gamma)$ is also a pseudogroup presentation for $\mathcal{G}$. The map associated with an element $w \in \bar{\pi}_1(\Gamma)$ is the union of all maps $\alpha \in F(\Gamma)$ that project to w. We prove that this map is in fact well-defined with connected domain.

Adopting the notation of [Gaboriau et al. 1992], we say that a pseudogroup presentation Γ has *independent generators* if $F(\Gamma)$ is isomorphic to $\bar{\pi}_1(\Gamma)$. If Γ has independent generators and $F(\Gamma)$ contains no reflection, $F(\Gamma)$ contains no cycle, for the existence of a cycle would imply a nontrivial relation among the basis elements of a free group, which is absurd. Section 3 contains a new proof of Levitt's result on independent generators:

Theorem 3.10. [Levitt 1990b, Theorem 5] *Let Γ be a finite pseudogroup presentation generating a pseudogroup $\mathcal{G}$ defined on a tree Z. Suppose $F(\Gamma)$ contains no reflection. If $\bar{\pi}_1(\Gamma)$ is a free group, there exists a finite presentation Ψ of $\mathcal{G}$ that has independent generators.*

Recently, Gaboriau has significantly improved this result (in preparation).

Let Γ be a pseudogroup defined on a tree Z, and let I be a subinterval of Z containing more than one point. Then I is a *dead end* of Γ if it has maximal length with respect to the following property: Each $p \in I$ is in the domain of exactly one element of $\Gamma \cup \Gamma^{-1}$. One may *remove* this dead end from Γ by restricting the domain of the offending map. The orbits of the new pseudogroup Γ' are trivial on I. The orbits of Γ' contained in $Z \setminus I$ equal the orbits of Γ restricted to $Z \setminus I$. If Γ' has a dead end, one can again remove it. The process of removing dead ends may uncover dead ends of smaller and smaller length, and so continue forever. Such a pseudogroup presentation gives rise to a *Levitt pseudogroup*. I am grateful to F. Paulin for pointing this out to me.

A pseudogroup presentation is *open* if its generators and their inverses are defined on open intervals of Z containing no vertices of Z. An open pseudogroup presentation Γ with a finite number of generators defined on a compact tree is an *interval exchange* if all but a finite number of points of Z lie in the domain of exactly zero or two maps of $\Gamma \cup \Gamma^{-1}$: see [Veech 1978], for example. If the process of dead-end removal discussed above terminates, one obtains such an interval exchange. This result is due to Levitt. The proof in Section 4 is based on the point of view of [Morgan 1988].

Theorem 4.4. [Levitt 1990a, Corollary II.5] *Suppose Γ is an open pseudogroup presentation defined on a compact tree Z, having a finite number of generators, and such that $\bar{\pi}_1(\Gamma)$ is isomorphic to $F(\Gamma)$, the free group on the set Γ. If Γ has no dead ends, it is an interval exchange.*

Section 5 gives the details of the connection between **R**-trees and pseudogroups as I understood them as of July 1991, at the Isle of Thorns conference. See [Rimlinger a; b; c] for more results along these lines, and [Gaboriau et al. 1992] for a complete proof of Rips' theorem (see Section 1) based on the trichotomy implied in Theorem 5.8.

Theorem 5.8. (Rips: see Section 1) *Let $G \times G \to T$ be a free minimal action of a finitely presented group G on an **R**-tree T. Let $\mathcal{G}$ be a pseudogroup generated by an open pseudogroup presentation Γ corresponding to this action. Then either*

(i) *$G = G' * H$, where H is a noncyclic free abelian group,*

or there exists a finite presentation Ψ of $\mathcal{G}$ with independent generators such that either

(ii) *the dead ends can be removed from* Ψ *in a finite number of steps, and the corresponding open pseudogroup presentation is an interval exchange, or*

(iii) *the dead ends cannot be removed from* Ψ *in a finite number of steps.*

Indirectly, Rips [Morgan 1991] proves that case (ii) implies $G = G' * H$, where H is a surface group, and that case (iii) implies $G = G' * F$, where F is a free group of rank n for some $n \geq 3$.

The principal tool involved in the proof of Theorem 5.8 is the *fundamental group* of a pseudogroup $\mathcal{G}$, denoted by $\pi_1(\mathcal{G})$. Following [Haefliger 1984], Levitt considered $\pi_1(\mathcal{G})$ in the context of foliated manifolds and deduced that $\pi_1(\mathcal{G})$ is a free product of free abelian groups [Levitt 1990b]. In [Levitt 1990a], $\pi_1(\mathcal{G})$ was considered in the context of partial rotations defined on open subintervals of the circle. Levitt showed that the result of [Levitt 1990b] cited above holds in this somewhat more combinatorial setting, *mutatis mutandis*. Gusmao (in preparation) extended the result to the case of orientation-reversing maps defined on open subintervals of the circle. This extension is cited in my proof of Theorem 5.8. In [Gaboriau et al. 1992], a similar result is used to analyze Rips' "unidentified combinatorial objects".

For arbitrary pseudogroup presentations Γ, the group $\bar{\pi}_1(\Gamma)$ is not an invariant of the orbit space $\mathcal{G}$ generated by Γ (see Remark 2.5). For the case of an open pseudogroup presentation Γ generating $\mathcal{G}$, there is an exact relationship between $\bar{\pi}_1(\Gamma)$ and Levitt's fundamental group: First translate Γ into a presentation Γ' defined on the circle, and let $\mathcal{G}'$ be the pseudogroup generated by Γ'. Then $\pi_1(\mathcal{G}')$ is isomorphic to $\bar{\pi}_1(\Gamma) * \mathbf{Z}$ (see Remark 5.4). Hence $\bar{\pi}_1(\Gamma)$ is free if and only if $\pi_1(\mathcal{G})$ is.

I was led to consider the notion of Nielsen transformations described in this paper by working with MacRTree, a computer program I developed for the Macintosh II. Initially, I felt that there were lots of exotic finitely generated groups that acted freely on $\mathbf{R}$-trees, and that these groups could be hunted down by computer simulations. Instead, the computer obstinately refused to yield any potential counterexamples, and I began to wonder why. As I reread the first few pages of [Lyndon and Schupp 1987], I began to realize that the computer was "doing Nielsen transformations". In Section 6, I expand on this discussion of the specific role of the computer in the development of my proofs of Theorems 3.10, 4.4 and 5.8.

I am grateful to Gilbert Levitt for correspondence that led me to an understanding of his work and potential applications to $\mathbf{R}$-trees.

1. BACKGROUND

In this section, I attempt to clarify the relationship between my work and that of E. Rips and G. Levitt. At the Isle of Thorns conference in July 1991, Rips gave a major exposition of his result that every finitely generated group G that acts freely on an $\mathbf{R}$-tree is a free product of free abelian groups and surface groups [Morgan 1991]. His proof was based on a notion of combinatorial complexity that Razborov [1985] used to investigate equations in free groups.

I arrived at the conference planning to speak about a process that constructs independent generators via Nielsen transformations [Nielsen 1921] for a pseudogroup with free fundamental group. I further wished to indicate the implications of this work for the classification problem for free actions on $\mathbf{R}$-trees. (Levitt had previously developed a similar result [Levitt 1990b, Theorem 5] in the context of foliated manifolds. I will return to this point below.) I quickly learned that Rips' proof was now accepted by many as correct. I then spoke with F. Paulin and learned that the process of removing dead ends from a pseudogroup presentation might not converge. This is the so-called Levitt pseudogroup case, advanced in [Levitt 1990a] as an indication that a previously proposed proof of the classification theorem associated with Rips was incomplete.

Therefore, I realized that there existed two major gaps in my understanding of a proof of the classification theorem, namely, how to prove that my method of Nielsen transformations did not alter the action of the corresponding action on an $\mathbf{R}$-tree, and how to deal with the Levitt pseudogroup case.

The substance of my talk, hastily revised to account for my conversation with Paulin and the news about Rips' proof, was that by finding independent generators for the pseudogroup the natural trichotomy between the free abelian case, the interval exchange case and the Levitt pseudogroup case could be clearly illustrated. The principal tool I used in this talk was Levitt's classification of the fundamental group of an open pseudogroup as a free product of free abelian groups [Levitt 1990a], as extended by Gusmao (in preparation). But, at that time, I did not know exactly how to translate the trichotomy involving pseudogroups to the corresponding trichotomy for free actions: translation along $\mathbf{R}$, surface group action and nonsimplicial free action.

Rips' proof uses an inductive procedure that applies uniformly to all three cases. The third case, the Levitt pseudogroup case, involves subtle convergence problems. Without Rips' solution of this case, the classification problem might well have remained open indefinitely.

After the Isle of Thorns conference, Gaboriau, Levitt and Paulin produced a complete proof of Rips' theorem [Gaboriau et al. 1992], based on the trichotomy of cases induced by the explicit construction of independent generators for the fundamental group of the pseudogroup. Levitt's result [Levitt 1990b, Theorem 5] provided a point of departure for their construction of independent generators. The process expounded in [Gaboriau et al. 1992] produces a pseudogroup with independent generators such that the corresponding group acting freely on an $\mathbf{R}$-tree does not change. However, this method alters the underlying set on which the pseudogroup is defined. My technique of Nielsen transformations produces independent generators that identically reproduce the orbits of the original pseudogroup on the original set, but the number of independent generators may be larger than the original number of generators. Gaboriau has improved this result by showing that it is not necessary to increase the number of generators to obtain independent generators.

I can now prove that the Nielsen transformation of a pseudogroup presentation does not change the associated group that acts freely on an $\mathbf{R}$-tree [Rimlinger a; b]. In fact, any action of G based on a finite presentation of G will remain invariant, not just the group G. This improvement in the normalization process has led me to an analysis of defining relators in surface groups [Rimlinger c]. One interesting point that needs more research is whether the program MacRTree (Section 6) can be shown to converge on surface group relators, possibly after the introduction of some kind of perturbation technique. Such a result would lead to significant improvements in the results of [Rimlinger c].

2. THE REDUCED FUNDAMENTAL GROUP

Trees and Graphs

An $\mathbf{R}$-*tree* is a path-connected metric space where each arc is isometric to an interval of $\mathbf{R}$. A *tree* Z is a connected, simply connected one-complex with a metric that gives Z the structure of an $\mathbf{R}$-tree and induces the topology of the one-complex structure. If I and J are connected subsets of a tree Z, an isometry $\gamma : I \to J$ is called a *partial isometry* of Z. We emphasize the fact that a partial isometry has connected domain and image by definition. A *graph* is defined to be a connected, simplicial one-complex.

Free Groups

The free group on a set X is denoted by $F(X)$. The set X^{-1} is defined as $X^{-1} = \{x^{-1} \in F(X) : x \in X\}$. Each nonidentity element $w \in F(X)$ can be uniquely expressed as a *reduced word* $w = w_n w_{n-1} \dots w_1$, with *letters* $w_i \in X \cup X^{-1}$, such that $w_{i+1} \neq w_i^{-1}$ for all i. We say n is the *length* of w, and denote it by $\ell_X(w)$, or simply $\ell(w)$. By convention, the identity element has length zero.

In general, we say $w = w_n w_{n-1} \dots w_1$ is reduced if $\ell(w) = \sum_{i=1}^n \ell(w_i)$. This allows for the appearance of the identity element at any point in a reduced word. (Of course, we actually mean to say that the ordered n-tuple $(w_n, w_{n-1}, \dots, w_1)$ is reduced, but this distinction must be inferred from context.) Now suppose w_1 and w_2 in $F(X)$ are such that the product $w_1 w_2$ has cancellation, that is, is not reduced. Then there exists a unique $u \neq 1$ and $a, b \in F(X)$ such that w_1, w_2 and $w_1 w_2$ may be expressed without cancellation as $w_1 = au^{-1}$, $w_2 = ub$ and $w_1 w_2 = ab$. In this event, we say that the parts u^{-1} of a and u of b have cancelled. In general, $\ell(w_1 w_2) = \ell(w_1) + \ell(w_2) - 2\ell(u) \leq \ell(w_1) + \ell(w_2)$.

Pseudogroups

Let Z denote an arbitrary fixed tree.

Definition 2.1. By a *pseudogroup presentation* Γ defined on Z we mean a set $\{\gamma : I \to J\}$ of partial isometries of Z.

Let Γ be a pseudogroup presentation defined on Z. Let $F(\Gamma)$ be the free group with basis Γ. For each element $\alpha \in F(\Gamma)$, we define a partial isometry $\Gamma(\alpha)$ as follows. Let $\Gamma(1)$ be the identity map on Z. If $\alpha \in F(\Gamma)$ has length 1, set $\Gamma(\alpha) = \alpha$ if $\alpha \in \Gamma$. If $\alpha \in \Gamma^{-1}$, set $\Gamma(\alpha) = \Gamma(\alpha^{-1})^{-1}$. Inductively, suppose $\alpha_n \alpha_{n-1} \dots \alpha_1 \in F(\Gamma)$ is reduced of length n, and $\Gamma(\alpha_n \alpha_{n-1} \dots \alpha_1)$ is defined with image I. Let $\alpha = \alpha_{n+1}\alpha_n \alpha_{n-1} \dots \alpha_1$ be reduced of length $n+1$. Let J be the domain of $\Gamma(\alpha_{n+1})$. Then $\Gamma(\alpha)$ is the composite $\Gamma(\alpha_{n+1}) \circ \Gamma(\alpha_n \alpha_{n-1} \dots \alpha_1)$ defined on $\Gamma(\alpha_n \dots \alpha_1)^{-1}(I \cap J)$. In particular, $\Gamma(\alpha)$ is the empty map if $I \cap J = \varnothing$. In this case, any reduced word that has a right segment α will also represent the empty map. Observe also that $\Gamma(\alpha\beta) = \Gamma(\alpha) \circ \Gamma(\beta)$ if $\alpha\beta$ is reduced. In general, $\Gamma(\alpha\beta)$ is an extension of $\Gamma(\alpha) \circ \Gamma(\beta)$.

We now dispense with the $\Gamma(\alpha)$ notation. Thus, given $p \in Z$ and $\alpha \in F(\Gamma)$, the expression $\alpha(p)$ implies that p is in the domain of $\Gamma(\alpha)$, and denotes the point $\Gamma(\alpha)(p) \in Z$.

Definition 2.2. We say that $p, q \in Z$ are in the same *orbit* generated by Γ if there exists $\alpha \in F(\Gamma)$ such that $q = \alpha(p)$. The set of equivalence classes $\mathcal{G}$ determined by this relation on Z is called the *pseudogroup* generated by Γ.

Definition 2.3. We say that a partial isometry α with domain $I \subset Z$ is a *reflection* if α has a fixed point $p \in I$, that is, if $\alpha(p) = p$ for some $p \in I$ and α is not the identity map of I.

Definition 2.4. Let Γ be a pseudogroup presentation. An element $\alpha \in F(\Gamma)$ such that $\alpha(p) = p$ for some $p \in Z$ is called a *cycle* of Γ. (Implicit in the statement $\alpha(p) = p$ is the assertion that α is indeed defined on p.) Let $N \subset F(\Gamma)$ be the normal subgroup of $F(\Gamma)$ generated by all the cycles of Γ. The quotient group $\bar{\pi}_1(\Gamma) = F(\Gamma)/N$ is the *reduced fundamental group* of Γ, and $\pi_\Gamma : F(\Gamma) \to \bar{\pi}_1(\Gamma)$ denotes the natural projection map.

Remark 2.5. Suppose Γ generates the pseudogroup $\mathcal{G}$. The reduced fundamental group $\bar{\pi}_1(\Gamma)$ is related to the fundamental group $\pi_1(\mathcal{G})$ defined in [Haefliger 1984] and [Levitt 1990a,b] (see Definition 5.3 below). However, if Ψ is another pseudogroup presentation that generates $\mathcal{G}$, it is not necessarily true that $\bar{\pi}_1(\Gamma)$ and $\bar{\pi}_1(\Psi)$ are isomorphic (compare Remark 5.4). For example, let Z be the closed interval $[0, 2]$ in $\mathbf{R}$. Let $\Gamma = \{\alpha, \beta\}$, where $\alpha(t) = t + 1$ for $0 \le t < \frac{1}{2}$, and $\beta(t) = t + 1$ for $1/2 \le t < 1$. Let $\Psi = \{\psi\}$, where $\psi(t) = t + 1$ for $0 \le t < 1$. Let $\mathcal{G} = \{\{t, t + 1\} : 0 \le t < 1\} \cup \{2\}$. Then both Γ and Ψ generate the pseudogroup $\mathcal{G}$. On the other hand, it is easy to see that both π_Γ and π_Ψ have trivial kernels, whence $\bar{\pi}_1(\Gamma)$ is a free group of rank two but $\bar{\pi}_1(\Psi)$ is a free group of rank one.

Each element of the reduced fundamental group of a pseudogroup presentation defined on Z can be regarded as a partial isometry of Z. We now pursue this line of thought, which will be needed in Section 3.

Lemma 2.6. *Let Γ be a pseudogroup presentation defined on the tree Z, and suppose $F(\Gamma)$ contains no reflection. Consider $\alpha, \beta \in F(\Gamma)$ such that there exists z in the domain of both α and β satisfying $\alpha(z) = \beta(z)$. Suppose $p, q \in Z$ are such that $\alpha(p)$ and $\beta(q)$ are defined. Then there is a partial isometry ψ defined on $[p, q]$ such that $\psi(t) = \alpha(t)$ if $\alpha(t)$ is defined, and $\psi(t) = \beta(t)$ if $\beta(t)$ is defined.*

Proof. Observe that z is a fixed point of the map $\beta^{-1}\alpha$. Since $F(\Gamma)$ contains no reflection, $\alpha(t) = \beta(t)$ for all $t \in \mathrm{domain}(\alpha) \cap \mathrm{domain}(\beta)$. Since Z is a tree, we know that the subtree of Z spanned by p, q and z has one of the three forms illustrated in Figure 1.

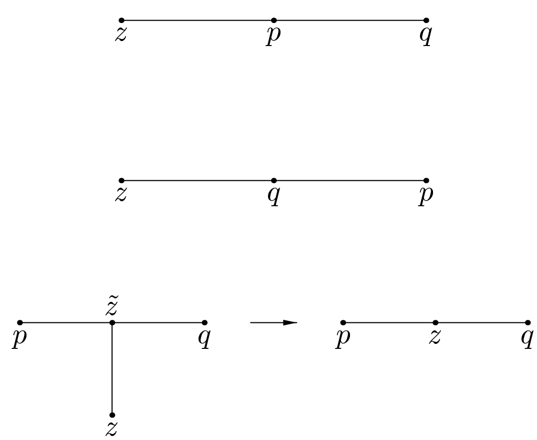

FIGURE 1. Possibilities for the subtree spanned by p, q and z.

By definition, the domains of α and β are connected. Thus α is defined on $[p, z]$ and β is defined on $[q, z]$. In case (i), define ψ to be the restriction of β to $[p, q]$. Likewise, in case (ii), ψ is the restriction of α to $[p, q]$. Notice that, in case (iii), the point labeled $\tilde{z}$ is also in the domains of α and β, so that we assume without loss of generality that

$z \in [p,q]$. Hence $d(p,q) = d(p,z) + d(z,q)$. Observe that $\alpha(z)$ and $\beta(z)$ are defined and equal, and denote this common value by z'. We know that $d(\alpha(p), z') = d(p,z)$ and $d(z', \beta(q)) = d(z,q)$. Thus $d(\alpha(p), \beta(q)) \leq d(p,z) + d(z,q)$. However, if this inequality were strict, we would deduce the existence of s in the images of both α and β satisfying $d(\alpha(p), s) < d(\alpha(p), z'))$ and $d(s, \beta(q)) < d(z', \beta(q))$ (see Figure 2). Since $\alpha^{-1}(s) = \beta^{-1}(s)$,

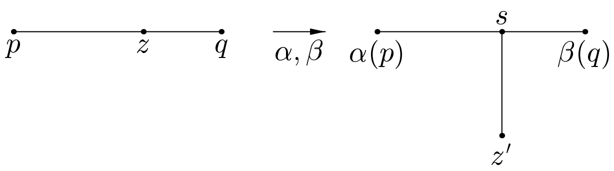

FIGURE 2. Illustrating the hypothetical case when $d(\alpha(p), \beta(q)) < d(p,z) + d(z,q)$.

we have the contradiction

$$\begin{aligned} d(p,q) &\leq d(p, \alpha^{-1}(s)) + d(\beta^{-1}(s), q) \\ &< d(p,z) + d(z,q) = d(p,q). \end{aligned}$$

Therefore $d(\alpha(p), \beta(q)) = d(p,z) + d(z,q)$, so that $z' \in [\alpha(p), \beta(q)]$. Defining $\psi(t) = \alpha(t)$ for $t \in [p,z]$ and $\psi(t) = \beta(t)$ for $t \in [z,q]$ yields the desired partial isometry of Z. □

The proof of the next lemma uses the notion of diagrams over the free group $F(\Gamma)$; see [Lyndon and Schupp 1987, p. 236] for the definition of a diagram. The only result we need about diagrams is the fact that they exist [Lyndon and Schupp 1987, Theorem 1.1]. The construction here differs from that of [Lyndon and Schupp 1987] to reflect the fact that elements of $F(\Gamma)$ are maps that compose naturally from right to left; this modification does not affect in any essential way the proof given there.

Lemma 2.7. *Let Γ be a pseudogroup presentation defined on the tree Z. Suppose $F(\Gamma)$ contains no reflection. Suppose that $\alpha, \beta \in F(\Gamma)$ are such that $\pi_\Gamma(\alpha) = \pi_\Gamma(\beta)$. Suppose $p, q \in Z$ are such that $\alpha(p)$ and $\beta(q)$ are defined. Then there exists a partial isometry $\psi : [p,q] \to [\alpha(p), \beta(q)]$ of Z such that $\psi(p) = \alpha(p)$, $\psi(q) = \beta(q)$ and for all $t \in [p,q]$, there exists $\theta_t \in F(\Gamma)$ such that $\theta_t(t) = \psi(t)$ and $\pi_\Gamma(\theta_t) = \pi_\Gamma(\alpha)$.*

Proof. Clearly $\beta^{-1}\alpha$ is in the kernel of π_Γ. We may assume that $\beta^{-1}\alpha$ is a finite product of conjugates of one or more cycles of Γ, say

$$\beta^{-1}\alpha = u_1 c_1 u_1^{-1} \dots u_n c_n u_n^{-1},$$

where each $c_i \in F(\Gamma)$ is a cyclically reduced cycle of Γ. By [Lyndon and Schupp 1987, Theorem 1.1], we may construct the Cayley diagram $M = M(c_1, \dots, c_n)$, a connected and simply connected planar two-complex. The two-cells of M are in one-to-one correspondence with the cycles c_i. We denote the one-skeleton of M by M^1. Each oriented edge of M^1 corresponds to an element of $F(\Gamma)$. Distinct edges with the same initial vertex correspond to elements γ_1 and γ_2 of $F(\Gamma)$ such that the product $\gamma_2^{-1}\gamma_1$ is reduced. An oriented arc with edges $E_0, \dots, E_{k-1}$ read off in order from first to last corresponds to the reduced word

$$\gamma_{k-1}\gamma_{k-2} \cdots \gamma_0$$

of $F(\Gamma)$, where each γ_i corresponds to E_i. Reading off the edges while traveling along the boundary ∂D of a two-cell D spells out, from right to left, a cyclically reduced product equal to the corresponding c_i or its inverse, up to cyclic permutation.

M has distinguished vertices V and W in its boundary. There are two oriented arcs from V to W in ∂M corresponding to α and β and whose union is equal to ∂M. In general, any path in M^1 with identical initial and terminal vertices corresponds to a word $\gamma \in F(\Gamma)$ whose projection $\pi_\Gamma(\gamma)$ is the trivial element of $\bar{\pi}_1(\Gamma)$.

Case 1. M is a single two-cell. Thus, $\beta^{-1}\alpha$ is cyclically reduced and equal to some cycle c of Γ. Take $z \in Z$ such that c is defined on z. Since c is an identity map, c^{-1} is also defined on z. Thus, α and β are both defined on z, and $\alpha(z) = \beta(z)$. By Lemma 2.6, there exists a partial isometry ψ defined on $[p,q]$ such that $\psi(p) = \alpha(p)$ and $\psi(q) = \beta(q)$. Setting $\theta_t(t) = \alpha(t)$ for $t \in [p,z]$ and $\theta_t(t) = \beta(t)$ for $t \in [z,q]$ proves Case 1.

Case 2. M is a union of n two-cells, and ∂M is a circle. Thus, $\beta^{-1}\alpha$ is cyclically reduced and equal to a product of conjugates of n cycles $c_i \in F(\Gamma)$. We take as an inductive hypothesis that if $\beta^{-1}\alpha$ can be expressed in this manner with fewer than n cycles, that is, if the number of two-cells of M can be reduced, the conclusions of the lemma hold.

For each vertex v in the boundary of a two-cell D of M, we define a point $z(v, D) \in Z$ as follows. Fix a cyclic ordering $v_0, \dots, v_{k-1}$ of the vertices of ∂D. For $i = 0, \dots, k-1$, let $\gamma_i \in F(\Gamma)$ correspond to the oriented edge $[v_i, v_{i+1}]$ of ∂D, with indices taken modulo k. Fix some $z \in Z$ defined on the cycle $\gamma_{k-1} \dots \gamma_0$ of Γ. Set $z(v_0, D) = z$, and set $z(v_i, D) = \gamma_{i-1} \dots \gamma_0(z)$ for $i = 1, \dots, k-1$.

For each oriented edge E of M with initial vertex $v \in M$ not equal to W, we define a closed interval $Z(E) \subset Z$ as follows. First suppose $E \subset \partial M$, so $E \subset \partial D$ for exactly one two-cell D of M. Let $\gamma \in F(\Gamma)$ correspond to the arc in ∂M from V to v that avoids W. Let $\mathcal{E} \subset \partial M$ be the arc from V to W containing E. If $\mathcal{E}$ corresponds to α, set $Z(E) = [\gamma(p), z(v, D)]$. If $\mathcal{E}$ represents β, set $Z(E) = [z(v, D), \gamma(q)]$. Finally, if $E \subset \partial D_1 \cap \partial D_2$ for distinct two-cells D_1 and D_2 of M, set $Z(E) = [z(v, D_1), z(v, D_2)]$.

The proof of Case 2 given below is based on the following simple observation. If $\gamma \in F(\Gamma)$ corresponds to an edge $E \subset M^1$, then γ is defined on both endpoints of $Z(E)$ and hence on the entire interval $Z(E)$.

Fix some $t \in [p, q]$. Define the graph $K_t \subset M^1$ as follows. Let $E_0, \dots, E_{k-1}$ be the oriented edges of an arc $\mathcal{E} = E_0 \cup \dots \cup E_{k-1} \subset M^1$ from V to some vertex $v \in M$. For $i = 0, \dots, k-1$, let $\gamma_i \in F(\Gamma)$ correspond to E_i. Say t *passes through* $\mathcal{E}$ if $t \in Z(E_0)$ and for $i = 1, \dots, k-1$, we have $\gamma_{i-1} \dots \gamma_0(t) \in Z(E_i)$. Let K_t be the union of $\mathcal{E}$ such that t passes through $\mathcal{E}$.

Proposition 2.8. *For all $t \in [p, q]$, either $W \in K_t$ or we can reduce the number of two-cells of M.*

For now we suppose the proposition is true, and complete the proof of Case 2. We assume that, $W \in K_t$ for all $t \in [p, q]$. This implies that, for all $t \in [p, q]$, there is an arc $\mathcal{E}_t \subset M^1$ from V to W such that t passes through $\mathcal{E}_t$. Fix $t \in [p, q]$, and let $\theta_t \in F(\Gamma)$ correspond to $\mathcal{E}_t$. Notice that $\mathcal{E}$, together with the arc from W to V in ∂M corresponding to α^{-1}, forms a path beginning and ending at V. Thus $\pi_\Gamma(\alpha^{-1}\theta_t)$ is the identity element of $\bar{\pi}_1(\Gamma)$, and clearly $\pi_\Gamma(\alpha) = \pi_\Gamma(\theta_t)$. Define $\psi : [p, q] \to Z$ by setting $\psi(t) = \theta_t(t)$, which is evidently defined. We verify that ψ has image $[\alpha(p), \beta(q)]$ and is a partial isometry of Z, as follows: Since there are only a finite number of paths in M from V to W, ψ may be constructed from a finite number of maps $\alpha = \theta_{t_1}, \dots, \theta_{t_k} = \beta$ such that each pair of maps $\theta_{t_i}, \theta_{t_{i+1}}$ have overlapping domains. Applying Lemma 2.6 to each such pair, we see that $\psi : [p, q] \to [\alpha(p), \beta(q)]$ is a partial isometry of Z.

Case 3. M consists of unions of two-cells bounded by circles joined together by arcs. This case follows immediately upon applying Case 2 to each union of two-cells of M whose boundary is a circle in ∂M, and then composing the maps yielded by these applications of Case 2 together with the maps along any connecting arcs. This concludes the proof of Lemma 2.7, assuming Proposition 2.8. □

Proof of Proposition 2.8. We use the following results, which we prove later.

Sublemma A. Let $\mathcal{E}_1$ and $\mathcal{E}_2$ be two oriented arcs in M^1 with initial vertex V and common terminal vertex v. Suppose $W \notin \mathcal{E}_1 \cup \mathcal{E}_2$, and suppose $t \in [p, q]$ passes through both $\mathcal{E}_1$ and $\mathcal{E}_2$. Let $\gamma_i \in F(\Gamma)$ correspond to $\mathcal{E}_i$ for $i = 1, 2$. Then $\gamma_1(t) = \gamma_2(t)$.

Sublemma B. V is an endpoint of an edge of K_t.

Sublemma C. Every vertex $v \in K_t$ distinct from V and W meets at least two distinct edges of K_t.

Sublemma D. Either $W \in K_t$, or the number of two-cells of M can be reduced, or K_t has the following property: If S is an embedded circle in K_t that bounds a two-cell D of M, either (a) $V \in S$, and there is an edge in $\overline{K_t \setminus S}$ that meets S, or (b) there exist two distinct edges in $\overline{K_t \setminus S}$ that meet S.

To prove the proposition, suppose first that K_t contains no embedded circle S, that is, K_t is contractible. By Sublemma B and the fact that K_t is compact, we see that K_t contains at least two distinct vertices $\{v, w\}$ such that $K_t \setminus \{v, w\}$ is connected. It now follows from Sublemma C that these two vertices must be V and W, so $W \in K_t$.

Now suppose that K_t does contain embedded circles. If $W \in K_t$, we are done; otherwise, Sublemma A implies that if an embedded circle bounds more than one two-cell, this circle corresponds to a cycle of Γ. Thus we may reduce the number of two-cells of M.

Now assume that each embedded circle in K_t bounds a single two-cell D of M. Choose an embedded circle S in K_t. Using Sublemma D, define the vertices v_1 and v_2 in S as follows. If $V \in S$, set $v_1 = V$ and let v_2 be the initial vertex of an edge of

$\overline{M^1 \setminus S}$ meeting S. Otherwise, let v_1 and v_2 be the (not necessarily distinct) vertices of two distinct edges E_1 and E_2 of $\overline{M^1 \setminus S}$ meeting S. It follows that distinct connected components of $S \setminus \{v_1, v_2\}$ lie within distinct components of $K_t \setminus \{v_1, v_2\}$ (otherwise, some circle of K_t would bound the union of D and at least one other two-cell of M, contradicting our assumption). Now let H be the graph obtained by removing from K_t a component of $K_t \setminus \{v_1, v_2\}$ that meets S and does not contain V. Observe that Sublemmas B–D remain valid when K_t is replaced by H. Since $\operatorname{rank} \pi_1(H) < \operatorname{rank} \pi_1(K_t)$, an induction on this rank allows to deduce that K_t contains a contractible subspace satisfying Sublemmas B–D. By the first case considered in the proof of this proposition, it now follows that $W \in K_t$.

Proof of Sublemma A. Suppose first that $\mathcal{E}_1 \cup \mathcal{E}_2$ is a circle, hence the boundary of a union of two-cells of M. It follows that $\pi_\Gamma(\gamma_2) = \pi_\Gamma(\gamma_1)$. Since $W \notin \mathcal{E}_1 \cup \mathcal{E}_2$, we see that $\mathcal{E}_1 \cup \mathcal{E}_2$ bounds fewer than n two-cells. By the inductive hypothesis of Case 2 in the proof of Lemma 2.7, the lemma holds for $\alpha = \gamma_1$, $\beta = \gamma_2$, and $p = t = q$. In particular, $\gamma_1(t) = \gamma_2(t)$, as desired. In general, $\mathcal{E}_1 \cup \mathcal{E}_2$ is a graph. The result follows upon applying the above argument to the outermost circles of this graph and composing the resulting maps with the maps represented by connecting arcs.

Proof of Sublemma B. Let $E_0, \dots, E_{k-1}$ be the edges of M^1 with initial vertex V, read off in order passing from the edge $E_0 \subset \partial M$ contained in the arc from V to W corresponding to α to the other edge $E_{k-1} \subset \partial M$ contained in the arc corresponding to β. For $i = 0, \dots, k-2$, the edges E_i and E_{i+1} lie in the boundary of a common two-cell D of M, thus $z(V, D) \subset Z(E_i) \cap Z(E_{i+1})$. We conclude that $\bigcup_{i=0}^{k-1} Z(E_i)$ is connected, whence $[p,q] \subset \bigcup_{i=0}^{k-1} Z(E_i)$, since $p \in Z(E_0)$ and $q \in Z(E_{k-1})$. Since $t \in [p,q]$, we conclude that t is in some $Z(E_i)$, as desired.

Proof of Sublemma C. Certainly, at least two oriented edges of M^1 have initial vertex v. Since $v \in K_t$ and $v \neq V$, one such edge, say E_0, lies in K_t. Now enumerate the edges of M^1 with initial vertex v as $E_0, \dots, E_{k-1}$, where $k \geq 2$. Notice that $Z(E_i) \cap Z(E_{i+1}) \neq \varnothing$, for $i = 0, \dots, k-1$, with indices taken modulo k. Let $\gamma \in F(\Gamma)$ correspond to an arc $\mathcal{E}$ from V to v through which t passes. By Sublemma A, it follows that $\gamma(t) \in Z(E_0)$. Now observe that both endpoints of $Z(E_0)$ are contained in $\bigcup_{i=1}^{k-1} Z(E_i)$ and that this union is a connected set. We conclude that $\gamma(t) \in \bigcup_{i=1}^{k-1} Z(E_i)$, whence t passes through $\mathcal{E} \cup E_i$ for some $i = 1, \dots, k-1$.

Proof of Sublemma D. Let $S \subset K_t$ be a circle bounding a two-cell D of M. Suppose there are less than two edges in $\overline{M^1 \setminus S}$ meeting S. Then $\partial D \subset \partial D'$ for some two-cell D' surrounding D. Thus $\partial(D \cup D') = \overline{\partial D' \setminus \partial D}$ corresponds to a cycle of Γ, and we can reduce the number of two-cells of M.

Henceforth, assume that $W \notin K_t$ and that there are at least two edges in $\overline{M^1 \setminus S}$ meeting S. Let $\gamma \in F(\Gamma)$ correspond to an arc from V to $v_0 \in S$ through which t passes. In the event $V \in S$, choose $v_0 = V$ and set $\gamma = 1$. Define $t_0 = \gamma(t)$. In the event $\overline{K_t \setminus S}$ contains an edge with initial vertex v_0, choose such an edge and call it E_0. We can always do this if $v_0 \neq V$. If $v_0 = V$ and no such E_0 exists, add an oriented edge E_0 to M^1 that meets M at the initial vertex V of E_0. Define $Z(E_0) = [p,q]$. Since $t \in [p,q]$ by hypothesis, we consider E_0 to be an edge of K_t. This notational device allows us to avoid considering $v_0 = V$ as a special case. Observe that, in general, $t_0 \in Z(E_0)$.

Fix an orientation of S and enumerate the oriented edges of $\overline{M^1 \setminus S}$ with initial vertex in S as $E_0, \dots, E_{k-1}$, where $k \geq 2$. Assume that each oriented edge E_i has initial vertex $v_i \in S$, and the oriented interval $[v_i, v_{i+1}]$ corresponds to $\gamma_i \in F(\Gamma)$, $i = 0, \dots, k-1$, with indices taken modulo k. Observe that $\gamma_i = 1$ if $v_i = v_{i+1}$.

For $i = 1, \dots, k-1$, set $t_i = \gamma_i \dots \gamma_0(t_0)$. By Sublemma A, these t_i are defined, since $S \subset K_t$. We assume that the third alternative of Sublemma D fails. This means that $t_i \notin Z(E_i)$ for $i = 1, \dots, k-1$. There exist points $z_0, \dots, z_{k-1}$ in Z such that $Z(E_0) = [\gamma_{k-1}(z_{k-1}), z_0]$ and

$$Z(E_i) = [\gamma_{i-1}(z_{i-1}), z_i]$$

for $i = 1, \dots, k-1$. Thus, for $i = 1, \dots, k-1$, the condition $t_i \notin Z(E_i)$ implies that $\gamma_{i-1}(z_{i-1})$ and z_i lie in the same component of $Z \setminus \{t_i\}$. Moreover, $t_{k-1} \neq z_{k-1}$ implies $t_0 \neq \gamma_{k-1}(z_{k-1})$, and $t_1 \neq \gamma_0(z_0)$ implies $t_0 \neq z_0$. Since $t_0 \in Z(E_0)$, we see that $\gamma_{k-1}(z_{k-1})$ and z_0 are in different components of $Z \setminus \{t_0\}$.

Now choose $t' \in Z$ in the component of $Z \setminus \{t_0\}$ containing z_0 and satisfying

$$d(t', t_0) < \min\{d(\gamma_{i-1}(z_{i-1}), t_i), d(z_i, t_i)\}$$

for $i = 1, \dots, k-1$. Then $\gamma_{k-1} \dots \gamma_0(t')$ is defined and lies in the component of $Z \setminus \{t_0\}$ containing $\gamma_{k-1}(z_{k-1})$. Hence, $\gamma_{k-1} \dots \gamma_0$ is a reflection, which contradicts the fact that $F(\Gamma)$ contains no reflections. Thus the third alternative of Sublemma D must hold. This concludes the proof of the sublemma and of Proposition 2.8. $\square$

Lemma 2.7 enables us to regard each $w \in F(\Gamma)$ as a partial isometry. The domain I of w is the set of $p \in Z$ such that $\alpha(p)$ is defined and $\pi_\Gamma(\alpha) = w$ for some $\alpha \in F(\Gamma)$. By Lemma 2.7, I is connected. Given $p \in I$, define $w(p) = \alpha(p)$, where α is defined on p and $\pi_\Gamma(\alpha) = w$. Setting $p = q$ in Lemma 2.7, we deduce that w is a well-defined function. Observe also that Lemma 2.7 implies that $d(w(p), w(q)) = d(p, q)$ for $p, q \in I$, whence w is a partial isometry of Z. Thus, $\bar{\pi}_1(\Gamma)$ is a pseudogroup presentation, and clearly Γ and $\bar{\pi}_1(\Gamma)$ generate the same pseudogroup.

The following corollary indicates the relationship between the group and pseudogroup presentation structures of $F(\Gamma)$ and $\bar{\pi}_1(\Gamma)$. By convention, if $x, y \in \bar{\pi}_1(\Gamma)$, we denote by xy the element of $\bar{\pi}_1(\Gamma)$ obtained from x and y by group multiplication, and by $x \circ y$ the partial isometry of Z obtained by composition.

Corollary 2.9. *Let $\alpha \in F(\Gamma)$. Then $\pi_\Gamma(\alpha) \in \bar{\pi}_1(\Gamma)$ is an extension of α. Let $x, y \in \bar{\pi}_1(\Gamma)$. Then $xy \in \bar{\pi}_1(\Gamma)$ is an extension of the composite map $x \circ y$.*

The reader might be tempted at this point to think that $w \in \bar{\pi}_1(\Gamma)$ is the maximal extension of some $\alpha \in \Gamma$ such that $\pi_\Gamma(\alpha) = w$ with respect to the property that Γ and $\Gamma \cup \{\text{extension of } \alpha\}$ generate the same pseudogroup $\mathcal{G}$. In view of Remark 2.5, this statement is false, since it would imply that $\bar{\pi}_1(\Gamma)$ depends only on the orbit space $\mathcal{G}$.

3. NIELSEN TRANSFORMATIONS

The notion of a Nielsen transformation goes back to [Nielsen 1921]. We shall use the approach given in detail in [Lyndon and Schupp 1987, pp. 4–7]. Nielsen originally used his technique to reduce a set of words in a free group to a basis for the subgroup they generate. In our setting, Nielsen transformations are used to take the slack, so to speak, out of a pseudogroup presentation Γ generating a pseudogroup $\mathcal{G}$. At the outset, we greatly enlarge the number of partial isometries and then toss out unnecessary ones in an orderly manner. The result is a Nielsen-reduced pseudogroup Ψ that also generates $\mathcal{G}$. We then draw the important conclusion that $\pi_\Psi : F(\Psi) \to \bar{\pi}_1(\Psi)$ is a group isomorphism, that is, Ψ has *independent generators*.

We begin with a discussion of order inside a finitely generated free group $F(X)$. These ideas are taken directly from [Lyndon and Schupp 1987, pp. 4–7]. Fix a well-ordering of the elements of $X \cup X^{-1}$. This ordering induces a well-ordering of all of $F(X)$, namely, the lexicographical ordering of reduced-word representation. For example, if $X = \{x_1, x_2\}$, one may order $\{X \cup X^{-1}\}$ by the rules $x_i < x_i^{-1}$ and $x_i < x_j$ for $i < j$, in which case $x_1^{-1}x_2x_1 < x_1^{-1}x_2^{-1}$. The identity element of $F(X)$ is represented by the empty word, which by definition is the smallest word in the lexicographical order.

Definition 3.1. Let $F(X)^* = F(X)/{\sim}$, where $\sim$ is the equivalence relation generated by the rule $w \sim w^{-1}$ for all $w \in F(X)$. Denote the element $\{w, w^{-1}\}$ of $F(X)^*$ by w^*. We define a well-ordering $<$ of $F(X)^*$ as follows. Say the left half of a word w is the initial segment $L(w)$ of length m, where m is the greatest integer not exceeding $\frac{1}{2}(\ell(w) + 1)$. Given a word $w \in F(X)$, we may speak of

$$\min\{L(w), L(w^{-1})\} \in F(X).$$

The minimum is determined by the lexicographical ordering of $F(X)$ and depends only on w^*. If $w_1^*, w_2^* \in F^*(X)$, say that $w_1^* < w_2^*$ if and only if one of the following conditions holds:

(i) $\ell(w_1) < \ell(w_2)$;

(ii) $\ell(w_1) = \ell(w_2)$ and
$$\min\{L(w_1), L(w_1^{-1})\} < \min\{L(w_2), L(w_2^{-1})\};$$

(iii) $\ell(w_1) = \ell(w_2)$ and
$$\min\{L(w_1), L(w_1^{-1})\} = \min\{L(w_2), L(w_2^{-1})\},$$
$$\max\{L(w_1), L(w_1^{-1})\} < \max\{L(w_2), L(w_2^{-1})\}.$$

Observe that $<$ is a well-ordering of $F(X)^*$. We also define $\leq$ on $F(X)^*$ by setting $w_1^* \leq w_2^*$ if $w_1^* < w_2^*$ or $w_1^* = w_2^*$. The notation $w_1^* < w_2^*$ should not be confused with $w_1 < w_2$, which means that

w_1 is less than w_2, according to the lexicographical ordering of $F(X)$.

Remark 3.2. Let $p, q, c \in F(X)$. Suppose that pc and qc are reduced and $\ell(p) = \ell(q) \le \ell(c)$. A moment of thought reveals that $p < q$ (lexicographically) implies $(pc)^* < (qc)^*$.

For the rest of this section, we work with a fixed finite pseudogroup presentation Γ defined on a tree Z and generating the pseudogroup $\mathcal{G}$. Our standing hypothesis is that $F(\Gamma)$ contains no reflection and $\bar{\pi}_1(\Gamma)$ is free. Fix a basis X for $\bar{\pi}_1(\Gamma) = F(X)$ and an ordering of $X \cup X^{-1}$. Order $\bar{\pi}_1(\Gamma)$ lexicographically as above. Finally, well-order $\bar{\pi}_1(\Gamma)^*$ as in Definition 3.1.

Recall from Section 2 that each $\alpha \in F(\Gamma)$ is a partial isometry of Z, and $\pi_\Gamma(\alpha) \in \bar{\pi}_1(\Gamma)$ is an extension of α. Moreover, $\bar{\pi}_1(\Gamma)$ generates $\mathcal{G}$ as a pseudogroup presentation. We now consider the more general situation of pseudogroup presentations Ψ that are *dominated* by Γ.

Definition 3.3. Let Ψ be a pseudogroup presentation of $\mathcal{G}$, and $\pi : F(\Psi) \to \bar{\pi}_1(\Gamma)$ a group homomorphism. Suppose that, for each $\alpha \in F(\Psi)$, the partial isometry $\pi(\alpha) \in \bar{\pi}_1(\Gamma)$ is an extension of α. Then we say that $\pi : F(\Psi) \to \bar{\pi}_1(\Gamma)$ is *dominated* by Γ, or simply that Ψ is dominated by Γ.

For example, Corollary 2.9 implies that Γ dominates itself: $\pi_\Gamma(\alpha)$ is an extension of α for each $\alpha \in F(\Gamma)$.

We say that $\alpha \in F(\Psi)$ is *defined* if its domain is nonempty. If $\pi : F(\Psi) \to \bar{\pi}_1(\Gamma)$ is dominated by Γ, we refer to the length of the word $\pi(\alpha) \in \bar{\pi}_1(\Gamma)$ simply as $\ell_X(\alpha)$ instead of the more cumbersome $\ell_X(\pi(\alpha))$.

The next definition can be compared with the one in [Lyndon and Schupp 1987, p. 6].

Definition 3.4. Suppose Ψ is dominated by Γ. We say that Ψ is *Nielsen-reduced* if the following conditions are satisfied.

(N0) For all $\alpha, \beta \in \Psi$, and for all $p \in Z$, if α is defined on p, then $\alpha(p) \ne p$, and if $\alpha\beta$ is defined on p, then $\alpha\beta(p) \ne p$.

(N1) For all $\alpha, \beta \in \Psi \cup \Psi^{-1}$, if $\alpha\beta$ is defined and $\pi(\alpha\beta) \ne 1$, then $\ell_X(\alpha\beta) \ge \max\{\ell_X(\alpha), \ell_X(\beta)\}$;

(N2) for all $\alpha, \beta, \gamma \in \Psi \cup \Psi^{-1}$, if $\alpha\beta$, $\beta\gamma$, and $\alpha\beta\gamma$ are defined, and $\pi(\alpha\beta) \ne 1$ and $\pi(\beta\gamma) \ne 1$, then $\ell_X(\alpha\beta\gamma) > \ell_X(\alpha) - \ell_X(\beta) + \ell_X(\gamma)$.

Our goal is to find a Ψ that is dominated by Γ and Nielsen-reduced. We now construct a finite set of partial isometries Ψ_0 that constitutes a "first approximation" to such a Φ.

Definition 3.5. Let $x^* \in \bar{\pi}_1(\Gamma)^*$ be the maximum of the finite set of elements

$$\{\pi(\gamma)^* \in \bar{\pi}_1(\Gamma)^* : \gamma \in \Gamma\}.$$

Let $\Psi_0 \subset \bar{\pi}_1(\Gamma)$ be such that $\Psi_0 \cap \Psi_0^{-1} = \varnothing$ and

$$\Psi_0 \cup \Psi_0^{-1} = \{w \in \bar{\pi}_1(\Gamma) : w^* \le x^*, w \ne 1\}.$$

By definition, $\pi_\Gamma(\Gamma) \setminus \{1\} \subset \Psi_0 \subset \bar{\pi}_1(\Gamma)$, whence Ψ_0 generates $\mathcal{G}$. Observe that the inclusion map $\Psi_0 \subset \bar{\pi}_1(\Gamma)$ induces a group homomorphism $\pi : F(\Psi_0) \to \bar{\pi}_1(\Gamma)$. Endow the free group $F(\Psi_0)$ with the pseudogroup presentation structure induced by Ψ_0. By Corollary 2.9, we see that $\pi(\alpha) \in \bar{\pi}_1(\Gamma)$ is an extension of α for all $\alpha \in F(\Psi_0)$. Thus $\pi : F(\Psi_0) \to \bar{\pi}_1(\Gamma)$ is dominated by Γ.

Let $w \in \Psi_0$ and suppose that $w(p)$ is defined for some $p \in Z$. If $w(p) = p$, we have $\pi_\Gamma(\alpha) = w$ and $\alpha(p) = p$ for some $\alpha \in F(\Gamma)$. Thus α is a cycle of Γ, so $w = \pi_\Gamma(\alpha) = 1$, contradicting the definition of Ψ_0. Thus $w(p) \ne p$. Similarly, $x, y \in \Psi_0$ and $xy(p) = p$ implies $x = y^{-1}$, contradicting $\Psi_0 \cap \Psi_0^{-1} = \varnothing$. Thus $xy(p) \ne p$. We conclude that Ψ_0 satisfies condition (N0) of Definition 3.4. In general, Ψ_0 does not satisfy (N1) or (N2).

Definition 3.6. Suppose Ψ is dominated by Γ and satisfies (N0). Let (α, β) be a pair of elements of $\Psi \cup \Psi^{-1}$ such that $\alpha\beta$ is defined. We say (α, β) is (N1)-*reducible* if $\pi(\alpha\beta) \ne 1$ and

$$\ell_X(\alpha\beta) < \max\{\ell_X(\alpha), \ell_X(\beta)\}.$$

We say (α, β) is (N2)-*reducible* if $\ell_X(\alpha) = \ell_X(\alpha\beta)$ and $\pi(\alpha\beta)^* < \pi(\alpha)^*$.

Observe that the (N1)-reducibility of (α, β) implies that of $(\beta^{-1}, \alpha^{-1})$. In general, this is not true of (N2)-reducibility.

In the following remarks, we suppose that $\alpha, \beta \in \Psi \cup \Psi^{-1}$ and that Ψ is dominated by Γ. Set $x = \pi(\alpha)$ and $y = \pi(\beta)$.

Remarks 3.7. (a) Suppose $\alpha\beta \in F(\Gamma)$ is defined and $\pi(\alpha\beta) \ne 1$. Suppose (α, β) is (N1)-reduced, so that $\ell(xy) \ge \max\{\ell(x), \ell(y)\}$. Let u be the part of y that cancels into x. Then $\ell(x) \le \ell(xy) = \ell(x) + \ell(y) - 2\ell(u)$, so that $\ell(u) \le \ell(y)/2$. Similarly,

$\ell(y) \leq \ell(xy)$ implies $\ell(u) \leq \ell(x)/2$. Thus neither x nor y cancels more than halfway into xy.

(b) Suppose (α, β) is (N1)- or (N2)-reducible. Let $u \in \bar{\pi}_1(\Gamma)$ be such that

$$u^* = \max\{\pi(\alpha)^*, \pi(\beta)^*\}.$$

It follows immediately that $\pi(\alpha\beta)^* < u^*$, whence $\pi(\alpha\beta) \notin \{u, u^{-1}\}$.

(c) If (α, β) is (N1)-reducible and $x = 1$, then $\ell(y) < \max\{\ell(y), 0\} = \ell(y)$, which is a contradiction. Thus $x \neq 1$, and similarly, $y \neq 1$. If (α, β) is (N2)-reducible, we again have $y \neq 1$ (since $y = 1$ would imply $x^* < x^*$) and $x \neq 1$ (since $x = 1$ would imply $\ell(y) = 0$ and $y = 1$).

(d) Suppose (α, β) is (N1)-reducible. We show that $\beta \notin \{\alpha, \alpha^{-1}\}$. Certainly $\beta \neq \alpha^{-1}$, since we have $\pi(\alpha\beta) \neq 1$. Now suppose $\beta = \alpha$. Thus $\ell(x^2) < \ell(x)$. Suppose $x = au^{-1}$ and $x = ub$ are reduced and such that $x^2 = ab$ is reduced. Observe that $\ell(a) = \ell(b)$ and $\ell(a) + \ell(b) < \ell(u) + \ell(b)$, whence $\ell(a) = \ell(b) < \ell(u)$. From the formula $au^{-1} = ub$, with both sides reduced, we see that a is a left segment of u, and b is a right segment of u^{-1}. Since ab is reduced, we deduce that $a = b = 1$. Thus $x^2 = 1$, which implies $x = 1$, since the free group $\bar{\pi}_1(\Gamma)$ is *a fortiori* torsion-free. This contradicts the previous remark.

(e) Suppose (α, β) is (N2)-reducible. We claim that $\beta \notin \{\alpha, \alpha^{-1}\}$. First suppose $\beta = \alpha$. Suppose $x = au^{-1}$ and $x = ub$ are reduced and such that $x^2 = ab$ is reduced. Since $\ell_X(\alpha) = \ell_X(\alpha\beta)$ by hypothesis, we have $\ell(x) = \ell(x^2)$, so $\ell(a) + \ell(u^{-1}) = \ell(a) + \ell(b)$, whence $\ell(u^{-1}) = \ell(b)$. From the formula $au^{-1} = ub$, with both sides reduced, we infer that $a = u$. Since au^{-1} is reduced, we deduce that $a = u = 1$. Thus $x = au^{-1} = 1$, which contradicts Remark 3.7(c). Now suppose $\beta = \alpha^{-1}$. Since $\ell_X(\alpha) = \ell_X(\alpha\beta)$, we have $\ell(x) = \ell(xx^{-1}) = 0$, whence $x = 1$, contradicting Remark 3.7(c).

The next lemma should be compared with Proposition 2.2 in Chapter 1 of [Lyndon and Schupp 1987].

Lemma 3.8. *Suppose* Ψ *is dominated by* Γ *and satisfies* (N0). *Suppose no pair* (α, β) *of elements of* $\Psi \cup \Psi^{-1}$ *is* (N1)- *or* (N2)-*reducible. Then* Ψ *is Nielsen-reduced.*

Proof. It is clear that Ψ satisfies (N1). Let $\alpha, \beta \in \Psi \cup \Psi^{-1}$, and suppose that $\alpha\beta$, $\beta\gamma$ and $\alpha\beta\gamma$ are defined and that $\pi(\alpha\beta) \neq 1$ and $\pi(\beta\gamma) \neq 1$. Let $x = \pi(\alpha)$, $y = \pi(\beta)$, and $z = \pi(\gamma)$. By Remark 3.7(a), we have $x = ap^{-1}$, $y = pbq^{-1}$, $z = qc$, all reduced, and $xy = abq^{-1}$, $yz = pbc$, $xyz = abc$, all reduced. If $b \neq 1$,

$$\begin{aligned} \ell(xyz) &= \ell(x) - \ell(y) + \ell(z) + 2\ell(b) \\ &> \ell(x) - \ell(y) + \ell(z), \end{aligned}$$

and so (N2) holds for the triple (α, β, γ).

Suppose now that $b = 1$, that is, $x = ap^{-1}$, $y = pq^{-1}$, and $z = qc$. It follows that (N2) is violated, so we must show that this case is impossible. Observe first that $\ell(p) \leq \frac{1}{2}\ell(y)$, $\ell(q) \leq \frac{1}{2}\ell(y)$, and $\ell(p) + \ell(q) = \ell(y)$. It follows that $\ell(p) = \frac{1}{2}\ell(y) = \ell(q)$. Moreover, the part of x that cancels into xy is less than or equal to half of x, and likewise the part of z that cancels into yz is less than or equal to half of z. Thus $\ell(p) = \ell(q) \leq \min\{\ell(a), \ell(c)\}$.

Since $y \neq 1$ and $y = pq^{-1}$, either $p < q$ lexicographically, or else $q < p$. If $p < q$, Remark 3.2 implies that $(pc)^* < (qc)^*$, or, equivalently, that $(c^{-1}p^{-1})^* < (c^{-1}q^{-1})^*$. Thus

$$\begin{aligned} \pi(\gamma^{-1}\beta^{-1})^* &= (z^{-1}y^{-1})^* = (c^{-1}p^{-1})^* \\ &< (c^{-1}q^{-1})^* = (z^{-1})^* = \pi(\gamma^{-1})^*, \end{aligned}$$

which contradicts the fact that $(\gamma^{-1}, \beta^{-1})$ is (N2)-reduced. On the other hand, suppose $q < p$. By Remark 3.2, $(qa^{-1})^* < (pa^{-1})^*$. Thus $\pi(\alpha\beta)^* = xy^* = (aq^{-1})^* < (ap^{-1})^* = x^* = \pi(\alpha)^*$, which contradicts the fact that (α, β) is (N2)-reduced. Thus $b \neq 1$, so that Ψ satisfies (N2). ☐

Lemma 3.9. *There exists a pseudogroup presentation* Ψ *of* $\mathcal{G}$ *that is dominated by* Γ *and Nielsen-reduced.*

Proof. In Definition 3.5 we built a pseudogroup presentation Ψ_0 of $\mathcal{G}$ dominated by Γ and satisfying (N0). Let N be the finite number of partial isometries contained in Ψ_0. Let $\mathcal{H}(n)$ stand for the following inductive hypotheses:

(i) $\pi : \Psi_n \to \bar{\pi}_1(\Gamma)$ is dominated by Γ, and Ψ_n satisfies (N0).

(ii) Let $R_n = \{\pi(\alpha)^* \in \bar{\pi}_1(\Gamma)^* : (\alpha, \beta)$ or (β, α) is (N1)- or (N2)-reducible, with $\alpha, \beta \in \Psi_n \cup \Psi_n^{-1}\}$. Then the set R_n has at most $N - n$ elements.

(iii) If $R_n \neq \varnothing$, let x^* be the maximum of R_n. If $w \in \bar{\pi}_1(\Gamma) \setminus \{1\}$ is such that $w^* \leq x^*$, then w is in $\Psi_n \cup \Psi_n^{-1}$.

It is clear that Ψ_0 satisfies (i) and (ii) of $\mathcal{H}(0)$. If $R_0 = \varnothing$, Lemma 3.8 implies that Ψ_0 is Nielsen-reduced and we are done. Otherwise, $x^* \in R_0$ exists and Ψ_0 satisfies (iii).

Inductively, suppose Ψ_n satisfies $\mathcal{H}(n)$ and R_n is nonempty. By part (iii) of the inductive hypothesis, there exists $x \in \Psi_n \cup \Psi_n^{-1}$ such that for all $w \in \bar{\pi}_1(\Gamma)$, if $w^* \leq x^*$ and $w \neq 1$, then $w \in \Psi_n \cup \Psi_n^{-1}$. Choose the notation so that $x \in \Psi_n$. Let $\Lambda = \Psi_n \cup \Psi_n^{-1} \setminus \{x, x^{-1}\}$.

Sublemma A. Let $\gamma \in \Psi_n \cup \Psi_n^{-1}$ and set

$$(\alpha,\beta) \in \{(x,\gamma),(x^{-1},\gamma),(\gamma,x),(\gamma,x^{-1})\}.$$

Suppose (α,β) is (N1)- or (N2)-reducible. Then $\{\pi(\alpha\beta),\pi(\gamma)\} \subset \Lambda$.

Proof. Let $y = \pi(\gamma)$. Then $y^* \leq x^*$ by maximality of $x^* = \{x, x^{-1}\}$. By Remark 3.7(c), $y \neq 1$. Thus hypothesis (iii) implies that $y \in \Psi_n \cup \Psi_n^{-1}$. By Remarks 3.7(d) and (e), we see that $y \notin \{x, x^{-1}\}$. We conclude that $\pi(\gamma) = y \in \Lambda$ and $\pi(\alpha\beta) \neq 1$. By Remark 3.7(b), $\pi(\alpha\beta)^* < x^*$, whence $\pi(\alpha\beta) \in \Lambda$. This concludes the proof of Sublemma A.

Let $\gamma \in \Psi_n \cup \Psi_n^{-1}$. We say that a point p in the domain of x is a *point of reduction* if

(R1) (x,γ) is (N1)- or (N2)-reducible and $\gamma^{-1}(p)$ is defined, or
(R2) (x^{-1},γ) is (N1)- or (N2)-reducible and $\gamma^{-1}(x(p))$ is defined, or
(R3) (γ,x) is (N1)- or (N2)-reducible and $\gamma(x(p))$ is defined, or
(R4) (γ,x^{-1}) is (N1)- or (N2)-reducible and $\gamma(p)$ is defined.

Sublemma B. If p is a point of reduction, there exist maps $h,k \in \Lambda$ such that $h \circ k(p) = x(p)$.

Proof. Let $y = \pi(\gamma)$. From Sublemma A, $y \in \Lambda$. We deduce that $y^{-1} \in \Lambda$, since Λ is closed under inverses. We consider the four cases of a point of reduction. In case (R1), $xy \in \Lambda$ by Sublemma A. From part (i) of the induction hypothesis $\mathcal{H}(n)$ and from Corollary 2.9, $y^{-1}(p)$ is defined, and $xy \circ y^{-1}(p) = x \circ y \circ y^{-1}(p) = x(p)$. Thus, setting $h = xy$ and $k = y^{-1}$ yields the result. In case (R2), we set $h = y$, $k = y^{-1}x$. Notice that $k \in \Lambda$, since $x^{-1}y \in \Lambda$ by Sublemma A and from the fact that Λ is closed under inverses. In case (R3), we set $h = y^{-1}$, $k = yx$. Finally, in case (R4), we set $h = xy^{-1}$, $k = y$. This concludes the proof of Sublemma B.

Let $\hat{x}$ be the restriction of x to the points in its domain that are not points of reduction. We may regard $\hat{x}$ as a finite set of partial isometries, each defined on a connected component of the domain of x. Set $\Psi_{n+1} = \Psi_n \backslash \{x\} \cup \hat{x}$. Since $\Lambda \subset \Psi_{n+1} \cup \Psi_{n+1}^{-1}$, Sublemma B implies that Ψ_{n+1} generates $\mathcal{G}$.

Define a map $\pi : \Psi_{n+1} \to \bar{\pi}_1(\Gamma)$ as follows. For each $\gamma \in \Psi_n \backslash \{x\}$, the image $\pi(\gamma)$ is defined by part (i) of $\mathcal{H}(n)$, and $\pi(\gamma) = x$ for each $\gamma \in \hat{x}$. Clearly, $\pi(\gamma)$ is an extension of γ for all $\gamma \in \Psi_{n+1}$, so that $\pi : \Psi_{n+1} \to \bar{\pi}_1(\Gamma)$ is dominated by Γ. Evidently, $R_{n+1} \subset R_n \setminus \{x^*\}$, so that $|R_{n+1}| \leq N - (n+1)$, and Ψ_{n+1} satisfies parts (ii) and (iii) of $\mathcal{H}(n+1)$. By induction, we have $R_n = \varnothing$ for some $n \leq N$, and by Lemma 3.8, Ψ_n is Nielsen-reduced. $\square$

Theorem 3.10. [Levitt 1990b, Theorem 5] *Let Γ be a finite pseudogroup presentation generating a pseudogroup $\mathcal{G}$ defined on a tree Z. Suppose $F(\Gamma)$ contains no reflection. If $\bar{\pi}_1(\Gamma)$ is a free group, there exists a finite presentation Ψ of $\mathcal{G}$ that has independent generators.*

Proof. By Lemma 3.9, there exists a finite presentation Ψ of $\mathcal{G}$ and a group homomorphism $\pi : F(\Psi) \to \bar{\pi}_1(\Gamma)$ that is dominated by Γ and Nielsen-reduced. Suppose $p \in Z$, and $\alpha \in F(\Psi)$ is such that $\alpha \neq 1$ and $\alpha(p)$ is defined. We must prove that $\alpha(p) \neq p$. By Lemma 2.7, it suffices to show that $\pi(\alpha) \neq 1 \in \bar{\pi}_1(\Gamma)$.

Let $\alpha_n\alpha_{n-1}\dots\alpha_1 = \alpha$ be the reduced-word representation of $\alpha \in F(\Psi)$. Set $w_i = \pi(\alpha_i)$ for each $i = 1,\dots,n$. Since Ψ satisfies (N0), we deduce that $w_{i+1} \neq w_i^{-1}$ for each $i = 1,\dots,n-1$. Thus (N1) implies that $\ell(w_{i+1}w_i) \geq \max\{\ell(w_i),\ell(w_{i+1})\}$, and (N2) implies that each triple (w_{i+2},w_{i+1},w_i) satisfies $\ell(w_{i+2}w_{i+1}w_i) > \ell(w_{i+2}) - \ell(w_{i+1}) + \ell(w_i)$.

For $n < 3$, (N0) implies that $w \neq 1$. If $n = 3$, then (N1) and (N2) taken together imply that we can write $w_3 = ap^{-1}$, $w_2 = pbq^{-1}$, $w_1 = qc$, $w_3w_2 = abq^{-1}$, $w_2w_1 = pbc$ and $w_3w_2w_1 = abc$, with all right-hand sides reduced and $b \neq 1$ (see the proof of Lemma 3.8). In particular, $w \neq 1$.

Inductively, suppose for all words $w_n \dots w_1$ of length $n \geq 3$ with consecutive pairs satisfying (N1) and consecutive triples satisfying (N2), we have $w_n = ap^{-1}$, $w_{n-1} = pbq^{-1}$, $w_{n-2}\dots w_1 = qc$,

$$w_nw_{n-1} = abq^{-1},$$

$w_{n-1}w_{n-2}\dots w_1 = pbc$ and $w_n \dots w_1 = abc$, with all right-hand sides reduced. Now suppose w_{n+1} is such that $w_{n+1} = a_1p_1^{-1}$, $w_n = p_1b_1q_1^{-1}$, $w_{n-1} = q_1c_1$, $w_{n+1}w_n = a_1b_1q_1^{-1}$, $w_nw_{n-1} = p_1b_1c_1$ and $w_{n+1}w_nw_{n-1} = a_1b_1c_1$, with all right-hand sides reduced, and $b_1 \neq 1$. Evidently, q_1^{-1} is the part of w_n that cancels into w_{n-1}. It follows that $q_1 = p$.

Since $pc_1 = w_{n-1} = pbq^{-1}$, we deduce that $c_1 = bq^{-1}$, whence $w_nw_{n-1} = p_1b_1c_1 = p_1b_1bq^{-1}$, which is reduced. Thus

$$\begin{aligned}(w_nw_{n-1})(w_{n-2}\dots w_1) &= (p_1b_1bq^{-1})(qc)\\ &= p_1b_1bc,\end{aligned}$$

which is reduced. Now set $c_2 = bc$. Then we have $w_{n+1} = a_1p_1^{-1}$, $w_n = p_1b_1q_1^{-1}$, $w_{n-1}\dots w_1 = pbc = q_1c_2$, $w_{n+1}w_n = a_1b_1q_1^{-1}$, $w_nw_{n-1}\dots w_1 = p_1b_1c_2$ and $w_{n+1}w_n\dots w_1 = a_1b_1c_2$, with all right-hand sides reduced and $b_1 \neq 1$. By induction, $w \neq 1$ for all n. □

4. ERGODIC THEORY

In general, $\bar{\pi}_1(\Gamma)$ is closely related to groups that act freely on $\mathbf{R}$-trees. See [Rimlinger a; b] for further development of this point of view. These papers, as well as Theorem 5.8 below, are based on the pseudogroup theory of [Haefliger 1984] and [Levitt 1990a,b]. This theory requires the partial isometries to be defined on *open* intervals. Accordingly, we begin this section with a discussion of the relationship between general pseudogroup presentations and those whose maps are defined on open intervals. We then use a technique of Morgan [Morgan 1988] to prove Theorem 4.4, stated in the Introduction.

Let Γ be a pseudogroup presentation defined on the tree Z. An *endpoint* of a connected subset $I \subset Z$ is a point $p \in I$ such that no neighborhood of p contained in I is isometric to an open interval of $\mathbf{R}$. Let Z^0 be the vertices of Z. For each $\gamma \in \Gamma$, let $\mathring{\gamma}$ be the restriction of γ to the set J, defined as the domain of γ, minus its endpoints, minus $Z^0 \cup \gamma^{-1}Z^0$. Observe that

$$\operatorname{domain}(\mathring{\gamma}) \cup \operatorname{image}(\mathring{\gamma}) \subset Z \setminus Z^0.$$

We regard $\mathring{\gamma}$ as a pseudogroup presentation where each map is defined on a connected component of J. Following [Gaboriau et al. 1992], we set $\mathring{\Gamma} = \bigcup_{\gamma\in\Gamma} \mathring{\gamma}$. Thus $\mathring{\Gamma}$ is a pseudogroup presentation of Z such that no map of $\mathring{\Gamma} \cup \mathring{\Gamma}^{-1}$ has an endpoint in its domain, and each map of $\mathring{\Gamma} \cup \mathring{\Gamma}^{-1}$ has a domain contained in $Z \setminus Z^0$.

We say that a pseudogroup presentation Γ defined on a tree is *open* if $\Gamma = \mathring{\Gamma}$. Observe that if Γ is open, then so is Γ^{-1}. In general, if Γ is a pseudogroup presentation defined on a tree, there is a projection map $\mathring{\Gamma} \to \Gamma$ that sends each map in $\mathring{\gamma}$ to γ. This projection determines a group homomorphism $F(\mathring{\Gamma}) \to F(\Gamma)$. Clearly, every cycle of $\mathring{\Gamma}$ maps to a cycle of Γ, so there is an induced group homomorphism $\bar{\pi}_1(\mathring{\Gamma}) \to \bar{\pi}_1(\Gamma)$.

Let Γ be an open pseudogroup presentation generating the pseudogroup $\mathcal{G}$ defined on a compact tree Z. It follows that the number of components of $Z \setminus Z^0$ is finite. Fix an injective map $\kappa : Z\setminus Z^0 \to \mathbf{R}$ that is isometric on each component of $Z \setminus Z^0$. Define the pseudogroup Γ_κ on $\mathbf{R}$ as the collection of maps $\{\gamma_\kappa\}_{\gamma\in\Gamma}$, where $\operatorname{domain}(\gamma_\kappa) = \kappa(\operatorname{domain}(\gamma))$, $\operatorname{image}(\gamma_\kappa) = \kappa(\operatorname{image}(\gamma))$, and $\gamma_\kappa(t) = \kappa\gamma\kappa^{-1}(t)$ for all $t \in \operatorname{domain}(\gamma_\kappa)$. The correspondence $\gamma \to \gamma_\kappa$ induces a group isomorphism $\kappa : F(\Gamma) \to F(\Gamma_\kappa)$. It is easy to verify that $c \in F(\Gamma)$ is a cycle of Γ if and only if $c_\kappa \in F(\Gamma_\kappa)$ is a cycle of Γ_κ. It follows that $\gamma \to \gamma_\kappa$ induces a group isomorphism $\kappa : \bar{\pi}_1(\Gamma) \to \bar{\pi}_1(\Gamma_\kappa)$.

Recall that a pseudogroup presentation is *finite* if it contains a finite number of generators.

Remark 4.1. Let Γ be a finite open pseudogroup presentation defined on a compact tree Z, and let $p \in Z$. Set $\Gamma_0(p) = \{p\}$, and, inductively,

$$\Gamma_{n+1}(p) = \Gamma_n(p) \cup \{\gamma(r) \in Z : r \in \Gamma_n(p),\ \gamma \in \Gamma \cup \Gamma^{-1}\}.$$

We claim that the orbit of p generated by Γ has polynomial growth, that is, the sequence $|\Gamma_n(p)|$ is bounded above by $f(n)$ for some polynomial f. (In fact, this growth rate is an invariant of the pseudogroup generated by Γ: see [Levitt 1990a].) Fix $\kappa : Z\setminus Z^0 \to \mathbf{R}$ as above. Certainly, the growth rate of p with respect to Γ is the same as that of $\kappa(p)$ with respect to Γ_κ. Now extend each map of Γ_κ to all of $\mathbf{R}$ to obtain generators for a finitely generated subgroup of the isometry group of $\mathbf{R}$. Such groups are well known to have polynomial growth, hence the orbit of p has polynomial growth with respect to Γ.

Definition 4.2. By an *interval exchange* we understand an open pseudogroup presentation Γ with a finite number of generators, defined on a compact tree Z, such that any $p \in Z$, with finitely many exceptions, is contained in the domain of exactly zero or two elements of $\Gamma \cup \Gamma^{-1}$.

Definition 4.3. Let Γ be a pseudogroup defined on a tree Z, and let I be a subinterval of Z containing more than one point. Then I is a *dead end* of Γ if it has maximal length with respect to the following property: Any $p \in I$ is in the domain of exactly one element of $\Gamma \cup \Gamma^{-1}$.

Observe that, if Γ is an open pseudogroup and $p \in Z$ is contained in the domain of two elements $\alpha, \beta \in \Gamma \cup \Gamma^{-1}$, there exists a neighborhood of p isometric to an open interval of $\mathbf{R}$ with this property.

Theorem 4.4. [Levitt 1990a, Corollary II.5] *Suppose Γ is an open pseudogroup presentation defined on a compact tree Z, having a finite number of generators, and such that $\bar{\pi}_1(\Gamma)$ is isomorphic to $F(\Gamma)$, the free group on the set Γ. If Γ has no dead ends, it is an interval exchange.*

Proof. We adjust the proof of Theorem 7.1 in [Morgan 1988] to account for the pseudogroup presentation Γ. Morgan's original theorem was concerned with words whose letters alternated between two sides of a free product with amalgamation. In the present case, the words are just elements of $F(\Gamma)$.

Given any point $p \in Z$, we define the *multiplicity* $m(p)$ of p to be 1 less than the cardinality of

$$\{\gamma \in \Gamma \cup \Gamma^{-1} : \gamma(p) \text{ is defined}\}.$$

Remarks 4.5. (a) Suppose $\alpha(p)$ is defined for some $p \in Z$ and $\alpha \in \Gamma \cup \Gamma^{-1}$. Then $m(\alpha(p))$ is the number of reduced words of the form $\beta\alpha$ that are defined on p.

(b) Since each $\gamma \in \Gamma \cup \Gamma^{-1}$ has an open domain, we may assume that $m(p) \geq 2$ for all p in some nonempty open interval of Z. Otherwise Γ is clearly an interval exchange, since it has no dead ends and contains only a finite number of maps. Moreover, the hypothesis that Γ has no dead ends implies that $m(p) \geq 1$ for all p in the domain of some element of $\Gamma \cup \Gamma^{-1}$.

Let Y be the space of infinite sequences

$$(p, \alpha_1(p), \alpha_2\alpha_1(p), \ldots)$$

such that each $\alpha_n \ldots \alpha_1 \in F(\Gamma)$ is a reduced word defined on p. We topologize Y as a subspace of the infinite product $\prod_{i=0}^{\infty} Z$ and fiber Y over the disjoint union of the domains of the maps of $\Gamma \cup \Gamma^{-1}$. Specifically, let $B \subset Z \times \Gamma \cup \Gamma^{-1}$ be defined as

$$B = \{\text{domain}(\alpha) \times \{\alpha\} : \alpha \in \Gamma \cup \Gamma^{-1}\}.$$

The fiber map $Y \to B$ sends $(p, \alpha_1(p), \alpha_2\alpha_1(p), \ldots)$ to (p, α_1). This map is well-defined because the equality $\alpha_1(p) = \beta(p)$ for some $\beta \in \Gamma \cup \Gamma^{-1}$ implies that $\alpha_1\beta^{-1}$ is a relation in $\bar{\pi}_1(\Gamma)$. The hypothesis that $F(\Gamma)$ is isomorphic to $\bar{\pi}_1$ implies that the projection $F(\Gamma) \to \bar{\pi}_1(\Gamma)$ is an isomorphism, since finitely generated free groups are Hopfian. Thus $\beta = \alpha_1$.

Let $F(p, \alpha)$ denote the fiber over a point $(p, \alpha) \in B$. A basic open set of $F(p, \alpha)$, say

$$V(p, \alpha(p), p_2, \ldots, p_n),$$

for $n \geq 2$, consists of all sequences with an initial segment of the form

$$(p, p_1, p_2, \ldots, p_n) = (p, \alpha(p), \alpha_2\alpha(p), \ldots, \alpha_n \ldots \alpha_2\alpha(p)).$$

We define the *weight* of such a segment to be

$$\text{wt}(p, p_1, p_2, \ldots, p_n) = \frac{1}{m(p_1) \ldots m(p_{n-1})}.$$

Define a Borel measure $\mu(p, \alpha)$ on $F(p, \alpha)$ by setting

$$\mu(p, \alpha)V(p, p_1, p_2, \ldots, p_n) = \text{wt}(p, p_1, p_2, \ldots, p_n).$$

Thus, the measure of a basic open set is the weight of its initial segment. The axioms for a measure can be directly verified. By Remark 4.5(a), exactly $N = m(\alpha(p))$ elements $\{\beta_1, \ldots, \beta_N\}$ contained in $\Gamma \cup \Gamma^{-1}$ are such that $\beta_i\alpha(p)$ is defined and $\beta_i\alpha$ is reduced. Thus,

$$\begin{aligned}\mu(p, \alpha)F(p, \alpha) &= \sum_{i=1}^{N} \mu(p, \alpha)V(p, \alpha(p), \beta_i\alpha(p)) \\ &= \sum_{i=1}^{N} \frac{1}{m(\alpha(p))} = N\left(\frac{1}{N}\right) = 1.\end{aligned}$$

Thus, the total mass of $\mu(p, \alpha)$ is 1. The integral of the measures $\mu(p, \alpha)$ with respect to Lebesgue measure dp on B induced from the metric on Z gives a Borel measure μ on Y. It has total mass equal to the sum of the lengths of the domains of the elements of $\Gamma \cup \Gamma^{-1}$.

Let $S : Y \to Y$ be the truncation map

$$(p_0, p_1, p_2, \ldots) \to (p_1, p_2, \ldots).$$

We claim that μ is measure-preserving, that is, $\mu(U) = \mu(S^{-1}(U))$ for any measurable set $U \subset Y$. A basic open set U of Y has the form

$$U = \bigcup_{p \in I} \{V(p, \alpha(p), p_2, \ldots, p_n)\},$$

where each $p_i = \alpha_i \ldots \alpha_2\alpha(p)$, and $\alpha_n \ldots \alpha_2\alpha$ is reduced and defined for all p in an interval I. Additionally, we may assume that $m(p)$ is constant on I, say $N = m(p)$. Notice that N is the cardinality of the set

$$\{\beta \in \Gamma \cap \Gamma^{-1} : \alpha\beta \text{ is reduced, } \beta^{-1} \text{ is defined on } p\}.$$

Accordingly, let $\beta_1, \ldots, \beta_N$ be such that $\alpha\beta_i$ is defined on $q_i = \beta_i^{-1}(p)$. We compute that

$$\begin{aligned}
&\mu(S^{-1}(U)) \\
&\quad = \sum_{i=1}^{N} \int_{\beta_i{}^{-1}(I)} \mu(q_i, \beta_i) V(q_i, p, \alpha(p), p_2, \ldots, p_n)\, dq_i \\
&\quad = \sum_{i=1}^{N} \int_{\beta_i{}^{-1}(I)} \frac{1}{N} \left(\frac{1}{m(\alpha(p) m(p_2) \ldots m(p_{n-1})} \right) dq_i \\
&\quad = \frac{1}{N} \sum_{i=1}^{N} \int_{I} \mu(p, \alpha) V(p, \alpha(p), p_2, \ldots, p_n)\, dp \\
&\quad = \mu(U).
\end{aligned}$$

We now apply the following result, which can be found in [Morgan 1988, § 7.5].

Theorem 4.6 (First Ergodic Theorem). *Let (Y, μ) be a measure space of finite total measure, and let the map $S : (Y, \mu) \to (Y, \mu)$ be measure-preserving. Let $f : Y \to \mathbf{R}$ be integrable. Then*

$$\hat{f}(p) = \lim_{n \to \infty} \left(\frac{1}{n} \right) \sum_{k=0}^{n-1} f(S^k(p))$$

exists for almost all $p \in Y$. Furthermore, $\hat{f}$ is integrable and $\int_Y \hat{f} d\mu = \int_Y f d\mu$.

To use this theorem, define $f : Y \to \mathbf{R}$ by the formula

$$f(p_0, p_1, \ldots) = \ln(m(p_1)).$$

Now set

$$\begin{aligned}
f_n(p_0, p_1, \ldots) &= \sum_{k=0}^{n-1} f(S^k(p_0, p_1, \ldots)) \\
&= -\ln \operatorname{wt}(p_0, p_1, \ldots, p_{n+1}).
\end{aligned}$$

This paragraph and the next are quoted with slight alterations from [Morgan 1988, p. 621]. By the First Ergodic Theorem, $\hat{f} = \lim_{n \to \infty} f_n/n$ exists almost everywhere and has same integral as f. By Remark 4.5(b), $m(p) \geq 2$ for some nonempty open interval $(I, \alpha) \subset B$. Moreover, $m(p) \geq 1$ for almost all $(p, \beta) \in B$. Thus $\int_Y f d\mu = K\mu(B)$ for some $K > 0$. By Fubini's theorem, this means that there is a point $(p, \alpha) \in B$ such that $\hat{f}$ is defined almost everywhere on the fiber $F(p, \alpha)$, and $\int_{F(p,\alpha)} \hat{f} d\mu(p, \alpha) \geq K$. Hence, there is a subset $E \subset F(p, \alpha)$ of positive $\mu(p, \alpha)$-measure, say $\mu(p, \alpha)(E) = \nu > 0$, on which $\hat{f} \geq K$. There is an integer N such that for all $n \geq N$ and all $a \in E \subset F(p, \alpha)$, we have $f_n(a)/n > K/2$, or equivalently, $f_n(a) > nK/2$. By the formula for f_n, this means that for any $(p, \alpha(p), \ldots) \in E$, we have the inequality $-\ln \operatorname{wt}(p, \alpha(p), \ldots, p_{n+1}) > nK/2$, or, equivalently, that

$$\begin{aligned}
\mu(p, \alpha) V(p, \alpha(p), \ldots, p_{n+1}) &= \operatorname{wt}(p, \alpha(p), \ldots, p_{n+1}) \\
&< \exp(-nK/2)
\end{aligned}$$

for all $n \geq N$.

Since the total mass of E is ν and since sets of the form $V(p, \alpha(p), \ldots, p_{n+1}) \cap E$, as $(p, \alpha(p), \ldots, p_{n+1})$ ranges over initial sequences of points of E, form a covering of E, it follows that, for all $n \geq N$, there are at least $\nu \exp(nK/2)$ sequences of the form $(p, \alpha(p), \ldots, p_{n+1})$ that occur as initial sequences of points of E and *a fortiori* as points of $F(p, \alpha)$. This means that, for all $n \geq N$, the number of elements of $F(\Gamma)$ of length $n + 1$ defined at p is at least $\nu \exp(nK/2)$.

By Remark 4.1, the orbit of p generated by Γ has polynomial growth. Thus, there must be two distinct reduced words $\beta_1, \beta_2 \in F(\Gamma)$ such that $\beta_1(p) = \beta_2(p)$. As we noted earlier in this proof, the hypothesis that $F(\Gamma)$ and $\bar{\pi}_1(\Gamma)$ are isomorphic implies that the projection $F(\Gamma) \to \bar{\pi}_1(\Gamma)$ is an isomorphism. Thus $\beta_1 = \beta_2$, which is a contradiction; and therefore, the assumption made in Remark 4.5(b) cannot hold, that is, the case $m(p) \geq 2$ for all p in some nonempty open interval of Z is impossible. We conclude that Γ is an interval exchange. $\square$

5. FREE ACTIONS ON $\mathbb{R}$-TREES

A group G acts *freely* on an $\mathbf{R}$-tree T if for any $p \in T$, the only element $g \in G$ such that $gp = p$ is the identity. G acts *by isometries* if each $g : T \to T$ is an isometry. G acts *minimally* if there is no proper subtree $S \subsetneq T$ such that $GS = S$. For an arbitrary group G acting on metric spaces M and N, a continuous map $\varphi : M \to N$ is a *G-map* if $\varphi(gp) = g\varphi(p)$ for all $g \in G$ and $p \in M$. An action of G on a set M is *quasifree* if there exists a subgroup S of G such that, for all $p \in M$, the stabilizer of p by the action of G is S. If G acts quasifreely on M with stabilizer S, then G/S acts freely on M. If G_1 and G_2 act on M, their actions are *equivalent* if there is a group isomorphism $\iota : G_1 \to G_2$ such that $gp = \iota(g)p$ for all $g \in G_1$ and $p \in M$.

We assume in this section that $G \times T \to T$ is a free, minimal action by isometries of a finitely presented group G on an $\mathbf{R}$-tree T. By [Rimlinger 1992, Theorem 5.7], there is a compact graph K with universal cover τ and a $\pi_1(K)$-map $\varphi : \tau \to T$. (The space τ is a tree in the sense of this paper: τ inherits a metric structure from that of T. The map φ is a morphism in the category of folds of τ along an equivalence relation determined by the action of G on T. See [Rimlinger 1992] for details.) The action of $\pi_1(K)$ on τ is by covering translations. The action of $\pi_1(K)$ on T is quasifree, so that $\pi_1(K)/S$ acts freely on T, where S is the stabilizer of the action of $\pi_1(K)$. The groups $\pi_1(K)/S$ and G are isomorphic, and the actions of these groups on T are equivalent.

Let $b \in \tau$ be the basepoint of τ. Fix a presentation $\langle X : R\rangle$ for G, so that $R \subset F(X)$ is a finite set of relators and $G \approx F(X)/\langle\langle R\rangle\rangle$. Let $Z_\tau = \bigcup_{\alpha\in R}[b, \alpha b]$, so that Z_τ is a compact subtree of τ. Notice that $Z = \varphi Z_\tau$ is the continuous image of a compact connected set and hence a compact connected subset of T. It follows that Z has a simplicial structure that agrees with the metric of T. Thus Z is a compact tree.

Let $\Gamma_\tau \subset \pi_1(K)$ be the set of covering translations of τ that do not move Z_τ off of itself:

$$\Gamma_\tau = \{\alpha \in \pi_1(K) : \alpha Z_\tau \cap Z_\tau \neq \varnothing\}.$$

Observe that Γ_τ is a finite subset of $\pi_1(K)$. Since φ is a $\pi_1(K)$-map, we may view Γ_τ as a finite pseudogroup presentation defined on Z. Given $\alpha \in \Gamma_\tau$, the domain of α is $Z \cap \alpha^{-1}(Z)$. Let $\Gamma = \mathring{\Gamma}_\tau$, as defined at the beginning of Section 4. Thus Γ is an open pseudogroup with a finite number of generators defined on a compact tree Z.

Remark 5.1. From [Rimlinger 1992, Theorem 4.1], it follows that G is determined up to isomorphism by Γ. The proof of this relies heavily on ideas from [Rimlinger 1992], so we just give a sketch here. (We do not use this result deductively in this paper.) The map $\varphi : \tau \to T$ is determined by a certain equivalence relation on K, say $D \subset K \times K$. This equivalence relation is constructed in the proof of [Rimlinger 1992, Theorem 5.7]. However, Γ also determines an equivalence relation $D(\Gamma)$ on K as follows. Say $p, q \in K$ are equivalent if there are lifts $\tilde{p}, \tilde{q}$ to τ such that $\varphi(\tilde{p})$ and $\varphi(\tilde{q})$ are in Z, and some element of $F(\Gamma)$ takes $\varphi(\tilde{p})$ to $\varphi(\tilde{q})$. Let $D(\Gamma)$ be the induced *segment-closed* equivalence relation on K [Rimlinger 1992, Definition 2.1]. It is easy to see that $D(\Gamma) \subset D$. From [Rimlinger 1992, Theorem 4.1], there exists an $\mathbf{R}$-tree $\hat{T}$ and a commutative diagram of $\pi_1(K)$-maps,

$$\begin{tikzcd} \tau \arrow[r, "\hat{\varphi}"] \arrow[dr, "\varphi"'] & \hat{T} \arrow[d, dashed, "\Phi"] \\ & T \end{tikzcd}$$

where $\hat{\varphi} : \tau \to \hat{T}$ is induced by $D(\Gamma)$ and $\varphi : T \to \tau$ is induced by D. The action of $\pi_1(K)$ on $\hat{T}$ is quasifree. Let $\hat{S}$ be the stabilizer of the action of $\pi_1(K)$ on $\hat{T}$. The fact that $D(\Gamma)$ is constructed from a presentation for G is used to show that $\pi_1(K)/\hat{S} \approx \pi_1(K)/S \approx G$, although $\hat{T}$ and T may differ. The essential point is that $\hat{\varphi}(\alpha b) = \hat{\varphi}(b)$ for each $\alpha \in R$. Thus G is determined up to isomorphism by Γ.

Remark 5.2. Let $\rho_\tau : \Gamma_\tau \to G$ be the restriction of the map $\pi_1(K) \to \pi_1(K)/S \approx G$. Observe that ρ_τ induces a map $\rho_\tau : F(\Gamma_\tau) \to G$. Now suppose $\alpha \in \ker(F(\Gamma_\tau) \to \bar{\pi}_1(\Gamma_\tau))$, so that $\alpha(p) = p$ for some $p \in Z$. Since $\varphi : \tau \to T$ is a $\pi_1(K)$-map, we observe that $\rho(\alpha)$, viewed as an isometry of T, is an extension of α. Since G acts freely on T, we deduce that $\rho_\tau(\alpha) = 1$, so that $\rho_\tau : F(\Gamma_\tau) \to G$ induces a homomorphism $\rho_\tau : \bar{\pi}_1(\Gamma_\tau) \to G$. From Section 4, we have a group homomorphism $\bar{\pi}_1(\Gamma) \to \bar{\pi}_1(\Gamma_\tau)$. It follows that there is a group homomorphism $\rho : \bar{\pi}_1(\Gamma) \to G$ obtained by composing this homomorphism with ρ_τ.

An open pseudogroup presentation *defined on the circle* $\mathbf{R}/\mathbf{Z}$ is a finite collection Ψ of maps defined on open subintervals of $\mathbf{R}/\mathbf{Z}$ of the form $\psi(t) = a + t$ or $\psi(t) = a - t$, for some constant a. The collection of orbits generated by these maps and their inverses is the open pseudogroup *defined on the circle*. The maps of positive derivative are *orientation-preserving* and those of negative derivative are *orientation-reversing* [Levitt 1990a, § I.5]. We now pass from the open pseudogroup Γ defined on Z to an open pseudogroup Γ_κ defined on the circle. As in Section 4, choose an injection $\kappa : Z \setminus Z^0 \to \mathbf{R}$ that is locally isometric on each component of $Z \setminus Z^0$. Arrange for $\kappa(Z)$ to lie in the open interval $(0, 1) \subset \mathbf{R}$ by composing the original κ with a constant scaling factor followed by a translation. Now regard κ as a map from $Z \setminus Z^0$ to the circle $\mathbf{R}/\mathbf{Z}$. Observe that the image of κ avoids the image of 0 in $\mathbf{R}/\mathbf{Z}$. Define Γ_κ as in Section 4. It is easily verified that Γ_κ is an open pseudogroup presentation defined on the circle.

Definition 5.3. [Levitt 1990a, § I.5] Suppose Ψ is an open pseudogroup presentation defined on the circle $\mathbf{R}/\mathbf{Z}$, generating the open pseudogroup $\mathcal{G}$. The *fundamental group* $\pi_1(\mathcal{G})$ is defined relative to the universal covering map $\mathbf{R} \to \mathbf{R}/\mathbf{Z}$. For each $\psi : I \to J$ in Ψ, lift ψ to $\tilde{\psi} : \tilde{I} \to \tilde{J}$, where $\text{length}(\tilde{I}) = \text{length}(\tilde{J}) = \text{length}(I)$. Let $\tilde{\Psi} = \{\tilde{\psi}\}$, and let $\psi_0 : \mathbf{R} \to \mathbf{R}$ be the map $t \to t + 1$. Let $F(\tilde{\Psi} \cup \{\psi_0\})$ be the free group on the set $\tilde{\Psi} \cup \{\psi_0\}$. Then $\pi_1(\Gamma) = F(\tilde{\Psi} \cup \{\psi_0\})/\langle\langle N \rangle\rangle$, where N is the set of reduced words $\alpha \in F(\tilde{\Psi} \cup \{\psi_0\})$ such that $\alpha(t) = t$ for some $t \in \mathbf{R}$.

By [Haefliger 1984] and [Levitt 1990a], $\pi_1(\mathcal{G})$ depends only on $\mathcal{G}$, not on the specific presentation chosen to generate $\mathcal{G}$.

Remark 5.4. Let $\mathcal{G}_\kappa$ be the open pseudogroup generated by Γ_κ. We consider the relation between $\bar{\pi}_1(\Gamma)$ and $\pi_1(\mathcal{G}_\kappa)$. Notice that each $\gamma_\kappa : I \to J$ lifts to $\tilde{\gamma}_\kappa : \tilde{I} \to \tilde{J}$ such that $\tilde{I}$, $\tilde{J} \subset (0, 1)$. It follows that any reduced word in $F(\tilde{\Gamma}_\kappa \cup \{\psi_0\})$ that contains an occurrence of ψ_0 is not trivial in $\pi_1(\Gamma_\kappa)$. Thus $N \subset F(\tilde{\Gamma}_\kappa)$, where $F(\tilde{\Gamma}_\kappa)$ is the subgroup of $F(\tilde{\Gamma}_\kappa \cup \{\psi_0\})$ generated by $\tilde{\Gamma}_\kappa$. We deduce that $\pi_1(\mathcal{G}_\kappa) \approx F(\tilde{\Gamma}_\kappa)/\langle\langle N \rangle\rangle * \mathbf{Z}$. Clearly, the map $\gamma \to \gamma_\kappa \to \tilde{\gamma}_\kappa$ induces an isomorphism $\bar{\pi}_1(\Gamma) \approx F(\tilde{\Gamma}_\kappa)/\langle\langle N \rangle\rangle$, so that $\pi_1(\mathcal{G}_\kappa) \approx \bar{\pi}_1(\Gamma) * \mathbf{Z}$. Thus $\pi_1(\mathcal{G}_\kappa)$ is free if and only if $\bar{\pi}_1(\Gamma)$ is.

Definition 5.5. [Levitt 1990a, § I.4] A *complete minimal component* of an open pseudogroup $\mathcal{G}$ defined on $\mathbf{R}/\mathbf{Z}$ is a maximal open interval I such that

(i) for all $p \in I$, the orbit $\mathcal{G}_p$ of p is dense in I, and
(ii) for all $t \in \mathbf{R}$ and $y \in I$, if $[p, p+t]$ and $[q, q+t]$ are contained in I, then $q + t \in \mathcal{G}_{p+t}$ if and only if $q \in \mathcal{G}_p$.

We now state Levitt's structure theorem for open pseudogroups of the circle.

Theorem 5.6. (Levitt, Gusmao) *Let $\mathcal{G}$ be an open pseudogroup defined on $\mathbf{R}/\mathbf{Z}$. The fundamental group $\pi_1(\mathcal{G})$ is isomorphic to a free product of a finitely generated free group and a finite number of noncyclic finitely generated free abelian groups. These free abelian groups are in one-to-one correspondence with the complete minimal components of $\mathcal{G}$.*

This is proved in [Levitt 1990a, § I.5] for the case of orientation-preserving maps. Gusmao (in preparation) recently extended the result to the case of orientation-reversing maps.

Lemma 5.7. *If $\pi_1(\mathcal{G}_\kappa)$ contains a free abelian subgroup of rank greater than 1, so does G.*

Proof. Suppose $\pi_1(\mathcal{G}_\kappa)$ contains a noncyclic free abelian subgroup. By Theorem 5.6, $\mathcal{G}_\kappa$ has a complete minimal component I. Let $\mathcal{G}$ be the pseudogroup defined on Z generated by Γ. Given $p \in Z$, let $\mathcal{G}_p$ denote the orbit of p by the maps of Γ and their inverses. There exists an open interval $J \subset \kappa^{-1}(I)$ of Z and a point $p \in J$ such that

(i) $\mathcal{G}_p \cap J$ is dense in J, and
(ii) for all $t \in \mathbf{R}$ and $q \in J$, if $[p, p+t]$ and $[q, q+t]$ are contained in J, then $q + t \in \mathcal{G}_{p+t}$ if and only if $q \in \mathcal{G}_p$.

Now suppose $[p, p + t]$ and $[q, q + t]$ are contained in J, and $q \in \mathcal{G}_p$. By the compactness of $[p, p + t]$, there exists a finite sequence $\alpha_1, \ldots, \alpha_n$ in $\bar{\pi}_1(\Gamma)$ such that

$$[p, p + t] \subset \bigcup_{\alpha_i} \text{domain}(\alpha_i),$$

domain$(\alpha_i) \cap$ domain$(\alpha_{i+1}) \neq \varnothing$ for each i, and $\alpha_i(p+r) = q+r$ for every $r \in [0,t]$ such that $p+r \in$ domain(α_i). Each $\alpha_{i+1}^{-1}\alpha_i$ is trivial in $\bar{\pi}_1(\Gamma)$, being the image of a cycle in $F(\Gamma)$. Recalling the homomorphism $\rho : \bar{\pi}_1(\Gamma) \to G$ from Remark 5.2, we deduce that $\rho(\alpha_1) = \rho(\alpha_2) = \cdots = \rho(\alpha_n)$. Thus, there exists $g \in G$ such that $g[p, p+t] = [q, q+t]$, where g acts by translation in the amount $d(p,q)$.

Fix $p \in J$. By (i) and (ii) above, we may find q_1, $q_2 \in \mathcal{G}(p)$ and $t > 0$ such that

(iii) $d(p, q_1)$ and $d(p, q_2)$ are linearly independent,
(iv) $[p, p+t] \cup [q_1, q_1+t] \cup [q_2, q_2+t] \subset J$, and
(v) $d(p, q_1) + d(p, q_2) < t$.

Let $g_1, g_2 \in G$ be such that g_i translates $[p, p+t]$ to $[q_i, q_i+t]$, for $i = 1, 2$. It follows that $g_1g_2(p) = g_2g_1(p)$, so g_1 and g_2 commute in G. Since the translation lengths of g_1 and g_2 are linearly independent, we deduce that g_1 and g_2 generate a rank-two free abelian subgroup of G. $\square$

Theorem 5.8. (Rips: see Section 1) *Let $G \times G \to T$ be a free minimal action of a finitely presented group G on an $\mathbf{R}$-tree T. Let $\mathcal{G}$ be a pseudogroup generated by an open pseudogroup presentation Γ corresponding to this action. Then either*

(i) *$G = G' * H$, where H is a noncyclic free abelian group*,

or there exists a finite presentation Ψ of $\mathcal{G}$ with independent generators such that either

(ii) *the dead ends can be removed from Ψ in a finite number of steps, and the corresponding open pseudogroup presentation is an interval exchange, or*
(ii) *the dead ends cannot be removed from Ψ in a finite number of steps.*

Proof. Let $\mathcal{G}_\kappa$ be the pseudogroup presentation on $\mathbf{R}/\mathbf{Z}$ corresponding to Γ_κ. First suppose $\pi_1(\mathcal{G}_\kappa)$ is not free. By Theorem 5.6, $\pi_1(\mathcal{G}_\kappa)$ contains a noncyclic free abelian subgroup. By Lemma 5.7, G contains a noncyclic free abelian subgroup. By a theorem in [Morgan and Skora], $G = G' * H$, as desired.

Now suppose $\pi_1(\mathcal{G}_\kappa)$ is free. By Remark 5.4, $\bar{\pi}_1(\Gamma)$ is free. By Theorem 3.10, $\mathcal{G}$ is generated by a pseudogroup presentation Ψ such that $\bar{\pi}_1(\Psi)$ is isomorphic to $F(\Psi)$, the free group on the set Ψ. Let $\Psi = \{\psi_i\}$ and suppose $I \subset$ domain(ψ_i) is a dead end. Let $\hat{\psi}_i$ be the restriction of ψ_i to domain$(\psi_i) - I$. Thus $\Psi_1 = \{\psi_j : j \neq i\} \cup \{\hat{\gamma}_i\}$ is the pseudogroup presentation formed by removing the dead end I from Ψ. Since $\bar{\pi}_1(\Psi)$ is free, it follows that $\bar{\pi}_1(\Psi_1)$ is free. (In general, these two groups are not isomorphic: If domain$(\gamma_i) \setminus I$ is not connected, rank $\bar{\pi}_1(\Psi_1) >$ rank $\bar{\pi}_1(\Psi)$.) A similar argument applies for a dead end I in the image of ψ_i. Thus the process of dead-end removal preserves the freeness of the reduced fundamental group. Suppose the dead ends can be removed in a finite number of steps. We arrive at a pseudogroup presentation Ψ_n such that $\bar{\pi}_1(\Psi_n)$ is free. Thus $\bar{\pi}_1(\mathring{\Psi}_n)$ is free. Theorem 4.4 now implies that $\mathring{\Psi}_n$ is an interval exchange. $\square$

6. THE COMPUTER EXPERIMENT

The discovery by Morgan and Shalen that most surface groups act freely on $\mathbf{R}$-trees [Morgan and Shalen 1984; 1991] caused speculation about what other finitely generated groups, if any, could act freely on $\mathbf{R}$-trees. It was quickly realized that surface group actions corresponded to interval exchanges with a dense orbit. More generally, any finite pseudogroup presentation Γ defined on a compact tree Z such that $F(\Gamma)$ contains no reflection corresponds to a free action of a finitely generated group on an $\mathbf{R}$-tree (see [Rimlinger 1992, Theorem 4.1], for example).

Years ago, I realized that the technique used in [Morgan 1988] implied that if a pseudogroup presentation Γ with no dead ends was not an interval exchange, $F(\Gamma)$ must contain cycles (compare Theorem 4.4). (I mistakenly dismissed the case of presentations with dead ends as irrelevant.) I felt that such cycles corresponded to "spurious relations" in the corresponding group G acting freely on an $\mathbf{R}$-tree. I thought that the "exotic" free actions on $\mathbf{R}$-trees would be those in which these "spurious relations" exhibited some recognizable kind of pattern. In other words, if one fixed a point $p \in Z$ and displayed the finite set of points $\alpha(p) \in Z$ such that $\alpha \in F(\Gamma)$ has reduced length n and $\alpha(p)$ is defined, some kind of fractal-like image would emerge as n became large. I reasoned that perhaps one could infer from the stability properties of such an image that the corresponding group acting freely on an $\mathbf{R}$-tree was not finitely presented.

I therefore set out to "simulate free actions on $\mathbf{R}$-trees". By this, I meant plotting orbits in search of the elusive stable patterns of cycles. The first conceptual problem involved determining, for a given choice of Γ, whether or not $F(\Gamma)$ contained a reflection. I restricted my attention to the case where Z is a closed interval of $\mathbf{R}$, and empirically observed that if Γ contained an orientation-reversing map, $F(\Gamma)$ was sure to have a reflection. (Obviously there are surface-group presentations that contain orientation-reversing maps, but such presentations are not "stable" in some sense.) Therefore, I further restricted my attention to the case where Γ contains only orientation-preserving maps.

These early computer experiments were disappointing. No patterns emerged. I should have given up, but I could not let go of the idea that the exotic free actions were out there waiting to be discovered. Therefore, I decided to play the devil's advocate and tried to develop an algorithm that would either "simplify" every Γ to an interval exchange or discover a noncyclic free abelian subgroup of $\bar{\pi}_1(\Gamma)$. At that point I started working with the Macintosh and specifically with the Think C development system. The incredibly sophisticated programming environment of the Mac is not something dilettante programmers can readily embrace, but Think C gave me all the hints and support I needed to get started. My feeling was that my earlier experiments probably had failed because my code was faulty. I felt that the powerful real-time graphics available on the Mac would give me the positive feedback I needed to verify that algorithms were working as intended.

The algorithm for detecting interval exchanges was quite elementary. Starting with a presentation Γ_0 with no dead ends, there exist α and β in $\Gamma_0 \cup \Gamma_0^{-1}$ such that domain$(\alpha) = [p,q]$, domain$(\beta) = [p,r]$, and $[p,q] \subset [p,r]$. Therefore, one constructs Γ' from Γ_0 by replacing α with $\gamma = \alpha\beta^{-1}$, where domain$(\gamma) = \beta[p,q]$. In the implementation, Γ_0 actually has a somewhat special form determined by a "top word" and a "bottom word". This special form is inherited by Γ' and implies that $[p,q]$ is a dead end of Γ'. Discard this dead end, obtaining a presentation Γ_1. Now repeat the operation to get a sequence $\Gamma_0, \Gamma_1, \Gamma_2, \ldots$. If some Γ_k is an interval exchange, the pseudogroup $\mathcal{G}$ generated by Γ has a presentation consisting of Γ_k, together with a collection of maps corresponding to "removable dead ends" (Theorem 5.8). In the actual implementation, there may be several maps $\alpha_1, \ldots, \alpha_k$ defined on $[p,q]$. These maps are identifiable in the display because their domains and ranges are all drawn in the same color. All these maps are simultaneously replaced by $\beta^{-1}\alpha_1, \ldots, \beta^{-1}\alpha_k$. The interval $[p,q]$ then becomes a dead end and is discarded. I dubbed this algorithm the "relabeling algorithm" because of the way the colors dance on the screen as the new domains are drawn.

The algorithm for detecting noncyclic abelian subgroups of $\bar{\pi}_1(\Gamma)$ involved searching for elements of $F(\Gamma)$ with relatively small translation lengths defined on the same relatively large interval. This algorithm eventually evolved into Lemma 5.7. In the implementation, I quickly discovered that each time a noncyclic abelian subgroup was found, one could plot its rank against the detected subinterval of its complete minimal component (Definition 5.5). The resulting graphics are beautiful to watch—they look like skyscrapers sprouting up in a dense city skyline—but are unfortunately utterly devoid of any scientific meaning. I called this algorithm the "painting algorithm", in honor of the Macintosh system A-trap PaintRect.

To implement these algorithms, I overcame formidable technical challenges in the area of efficient dynamic memory management of data with a variable number of dimensions and a variable-type definition. I have developed intelligent general-purpose object-oriented memory management software, which I will be happy to share with other Think C developers.

And so MacRtree was born. I fed MacRtree dozens and dozens of presentations. In almost all cases, either the relabeling algorithm yielded an interval exchange after a finite number of steps, or the painting algorithm discovered a noncyclic abelian subgroup. In a few instances, the relabeling algorithm failed to converge when it "should" have. I attributed this to the "biased" way in which MacRtree chose the intervals $[p,q]$ and $[p,r]$. My feeling is that the relabeling algorithm can be "randomized" in such a way that it will always yield a positive result when fed a presentation generating a pseudogroup $\mathcal{G}$ such that $\mathcal{G}$ has an interval exchange presentation. Such a result would have implications for some recent theoretical work on presentations of surface groups [Rimlinger c].

At length, I realized that the relabeling algorithm was just a form of Nielsen transformation. From Lyndon and Schupp I gleaned the technical expertise to concoct a theoretical Nielsen transformation of Γ that converges to a pseudogroup Ψ with independent generators if $\bar{\pi}_1(\Gamma)$ is free (Theorem 3.10). From combining this result with the technique of [Morgan 1988], it follows that Ψ is an interval exchange if it has no dead ends (Theorem 4.4). To avoid worries about endpoints while performing the Nielsen reduction, I generalized the notion of Levitt's fundamental group to partial maps with connected domain (see Lemma 2.7 and Corollary 2.9). Considering these results in the light of Levitt's structure theorem for $\pi_1(\mathcal{G})$ (see Theorem 5.6), I began to believe there were no "exotic actions" living in the chimerical world between the interval exchange and the free abelian case (compare Theorem 5.8).

REFERENCES

[Alperin and Bass 1987] R. Alperin and H. Bass, "Length functions of group actions on Λ-trees", pp. 265–377, in *Combinatorial Group Theory and Topology* (edited by S. M. Gersten and J. R. Stallings), Ann. of Math. Studies **111** (1987).

[Chiswell 1976] I. M. Chiswell, "Abstract length functions in groups", *Math. Proc. Camb. Phil. Soc.* **80** (1976), 451–463.

[Gaboriau et al. 1992] D. Gaboriau, G. Levitt and F. Paulin, "Pseudogroups of isometries and Rips' theorem on free actions on **R**-trees", Rapport de Recherche 69, CNRS-UMR 128, Ecole Normale Supérieure de Lyon, 1992.

[Haefliger 1984] A. Haefliger, "Groupoïdes d'holonomie et classifiants", pp. 70–97, in *Structures Transverses des Feuilletages*, *Astérisque* **116** (1984).

[Levitt 1990a] Gilbert Levitt, *La Dynamique des Pseudogroupes de Rotations*, Prépublications du Laboratoire de Topologie et Géométrie URA CNRS **1408**, Université Paul Sabatier, Toulouse, 1990.

[Levitt 1990b] Gilbert Levitt, "Groupe fondamental de l'espace des feuilles dans les feuilletages sans holonomie", *J. Diff. Geom.* **31** (1990), 711–761.

[Lyndon 1963] Roger C. Lyndon, "Length functions in groups", *Math. Scand.* **12** (1963), 209–234.

[Lyndon and Schupp 1987] Roger C. Lyndon and Paul E. Schupp, *Combinatorial Group Theory*, Springer-Verlag, New York, 1987.

[Morgan 1988] J. Morgan, "Ergodic theory and free group actions on **R**-trees", *Invent. Math.* **94** (1988), 605–622.

[Morgan 1991] J. Morgan, "Notes on Rips' lectures at Columbia University", October 1991 (manuscript).

[Morgan and Shalen 1984] J. Morgan and P. Shalen, "Valuations, trees, and degenerations of hyperbolic structures I", *Ann. of Math.* **120** (1984), 401–476.

[Morgan and Skora] J. Morgan and R. Skora, "Groups acting freely on **R**-trees" (preprint).

[Morgan and Shalen 1991] J. Morgan and P. Shalen, "Free actions of surface groups on **R**-trees", *Topology* **30** (1991), 143–154.

[Nielsen 1921] J. Nielsen, "Om regning med ikke kommutative faktoren og dens anvendelse i gruppeteorien", *Mat. Tkdsskr.* **B** (1921), 77–94.

[Razborov 1985] A. A. Razborov, "On systems of equations in a free group", *Math. USSR Izvestiya* **25** (1985), 115–162.

[Rimlinger 1992] F. Rimlinger, "Free actions on **R**-trees", *Trans. Amer. Math. Soc.* **332** (1992), 315–331.

[Rimlinger a] F. Rimlinger, "Fibered disks and free actions on **R**-trees" (preprint).

[Rimlinger b] F. Rimlinger, "Two-complexes with similar foliations" (preprint).

[Rimlinger c] F. Rimlinger, "Untangling surface group relators" (preprint).

[Shalen 1987] P. B. Shalen, "Dendrology of groups: an introduction", in *Essays in Group Theory* (edited by S. M. Gersten), MSRI Publ. **8**, Springer-Verlag, New York, 1987.

[Veech 1978] W. A. Veech, "Interval exchange transformations", *Ann. of Math.* **33** (1978), 222–272.

Frank Rimlinger, Department of Mathematics, Fairfield University, Fairfield, CT 06430-7524
(frimling@fair1.fairfield.edu)

Received October 30, 1991; revised May 22, 1992