

THE FUNDAMENTAL GROUP OF A LOCALLY FINITE GRAPH WITH ENDS- A HYPERFINITE APPROACH

ISAAC GOLDBRING AND ALESSANDRO SISTO

ABSTRACT. The *end compactification* $|\Gamma|$ of the locally finite graph Γ is the union of the graph and its ends, endowed with a suitable topology. We show that $\pi_1(|\Gamma|)$ embeds into a nonstandard free group with hyperfinitely many generators, i.e. an ultraproduct of finitely generated free groups, and that the embedding we construct factors through an embedding into an inverse limit of free groups. We also show how to recover the standard description of $\pi_1(|\Gamma|)$ given in [DS11a]. Finally, we give some applications of our result, including a short proof that certain loops in $|\Gamma|$ are non-nullhomologous.

1. INTRODUCTION

It is a well-known fact that the fundamental group of a finite, connected graph Γ (viewed as a one-dimensional CW-complex) is a finitely generated free group. Indeed, fix a spanning tree T of Γ and let e_1, \dots, e_n denote the chords of T (that is, edges of Γ not in T) equipped with a fixed orientation. Then to any loop α , we can consider the word r_α on the alphabet $\{e_1^{\pm 1}, \dots, e_n^{\pm 1}\}$ obtained by recording the traversals of α on each e_i taking into account the orientation. If we let F_n denote the free group with basis e_1, \dots, e_n and we let $[r_\alpha]$ denote the unique reduced word corresponding to r_α , then the map $\alpha \mapsto [r_\alpha] : \pi_1(\Gamma) \rightarrow F_n$ is an isomorphism.

We now consider the case that Γ is an infinite, locally finite, connected graph. In order to obtain a compact space, we consider the *end compactification* $|\Gamma|$ of Γ obtained by adding the *ends of Γ* . Loosely speaking (we will make this more precise in the next section), the ends of Γ are the “path components of Γ at infinity”; see the paper [Gol11] where this heuristic is made precise using the language of nonstandard analysis. For example, if Γ is the Cayley graph of \mathbb{Z} with its usual generating set $\{1\}$, then Γ has two ends, one at “ $-\infty$ ” and one at “ $+\infty$.” End compactifications of graphs have been considered, e.g., in [Die10a, Die11, DK04, Geo09b, Geo09a, BB11] as a way to obtain analogues for infinite graphs of results in finite graph theory that would otherwise be plainly false.

Goldbring’s work was partially supported by NSF grant DMS-1007144.

The authors would like to thank the Centre International de Recontres Mathématiques in Luminy, France and the organizers of the meeting *Model Theory of Groups*, which took place in November 2011, where the authors first met to discuss this project.

In [DS11a], the authors find a combinatorial characterization of $\pi_1(|\Gamma|)$ similar in spirit to the aforementioned construction for finite graphs. The authors have to overcome two main obstructions, both of which are illustrated through the following simple example. Consider the 1-way infinite sideways ladder Γ indicated below in Figure 1.




FIGURE 1

Γ has one end, namely the point at “ $+\infty$.” The loop α that starts at v_0 , runs down the bottom edge of the ladder to $+\infty$ and then back to v_0 is clearly nullhomotopic. However, if one looks at the traces α leaves on the chords of the spanning tree T pictured in Figure 1, we are left with the following word:

$$r_\alpha = (e_1 e_2 \cdots)^\frown (\cdots e_2 e_1).$$

This word has order type $\omega + \omega^*$ (where ω is the order type of the natural numbers and ω^* is ω equipped with its reverse ordering) and contains no consecutive appearances of e_i, e_i^{-1} for any i . In particular, the (usual) notion of reduction of words does not work well in this context. More generally, for any infinite, locally finite graph Γ , the authors of [DS11a] consider the group F_∞ consisting of reduced infinite words on the alphabet of oriented chords of a *topological* spanning tree of Γ ; here, an infinite word can have arbitrary countable order type, leading to a complicated (non-well-ordered) notion of reduction of words. Just taking any spanning tree of Γ would not work, as the closure of such a tree in $|\Gamma|$ might contain loops. They then show that $\pi_1(|\Gamma|)$ embeds as a subgroup of F_∞ and F_∞ embeds as a subgroup of an inverse limit of finitely generated free groups by sending an infinite word to the family of its finite subwords. These are, respectively, part (i) and (ii) of [DS11a, Theorem 15].

After seeing the nonstandard approach to ends outlined in [Gol11], Diestel asked the first author whether or not the nonstandard approach could lead to a simplification of the article [DS11a]. Immediately the following idea arose: Consider the graph Γ from Figure 1 and fix an infinite natural number $\nu \in \mathbb{N}^* \setminus \mathbb{N}$. (In this paper, we will assume that the reader is familiar with the basics of nonstandard analysis; otherwise they may consult [Dav77] or [Hen97].) We then consider the *hyperfinite* graph Γ_ν appearing in Figure 2.

The nullhomotopic loop α from the previous paragraph “naturally” induces an *internal* loop in Γ^* , namely the loop that starts at v_0 , travels down the bottom edge to v_ν , then returns back to v_0 . The *hyperfinite* word that α corresponds to is $e_1 e_2 \cdots e_\nu e_\nu^{-1} \cdots e_2^{-1} e_1^{-1}$, which clearly reduces to the trivial word, witnessing that α is nullhomotopic.

$$\Gamma_\nu = \begin{array}{|c|c|c|c|c|c|} \hline & & & \cdots & & \\ \hline e_1 & e_2 & e_3 & \cdots & e_{\nu-1} & e_\nu \\ \hline v_0 & v_1 & v_2 & \cdots & v_{\nu-2} & v_{\nu-1} & v_\nu \\ \hline \end{array}$$

FIGURE 2

The main purpose of this article is to make precise the naïve approach taken in the previous paragraph. More specifically, we prove

Theorem 1.1. *Suppose Γ is a locally finite connected graph. Then there is a hyperfinite, internally connected graph Γ_{hyp} and an injective group morphism $\Theta : \pi_1(|\Gamma|) \hookrightarrow \pi_1(\Gamma_{\text{hyp}})$, where $\pi_1(\Gamma_{\text{hyp}})$ is the internal fundamental group of Γ_{hyp} . Moreover, Θ can be constructed in such a way that it factors through an embedding of $\pi_1(|\Gamma|)$ into an inverse limit of finitely generated free groups.*

By the transfer principle, $\pi_1(\Gamma_{\text{hyp}})$ will be an *internally free group on hyperfinitely many generators*. Said another way: $\pi_1(|\Gamma|)$ embeds into an ultraproduct of finitely generated free groups. The “moreover” part gives part of the aforementioned result of [DS11a]. In fact, in Section 4, we will show how our methods can be used to completely recover part (i) of [DS11a, Theorem 15] as described above (part (ii) is purely algebraic and the nonstandard set-up has nothing to say about this portion of the theorem). Two sample consequences of the theorem are reported below.

Corollary 1.2.

- (1) $\pi_1(|\Gamma|)$ has the same universal theory as a free group.
- (2) $\pi_1(|\Gamma|)$ is ω -residually free.

Recall that a group G is ω -residually free if for any $g_1, \dots, g_n \in G \setminus \{1\}$, there is a surjective homomorphism $\sigma : G \rightarrow F$ onto a free group such that $\sigma(g_i) \neq 1$ for $i = 1, \dots, n$. For finitely generated groups, conditions (1) and (2) in the above corollary are well-known to be equivalent, see e.g. [CG05]. Finitely generated groups satisfying conditions (1) and (2) are called *limit groups* and are widely studied.

We should mention that our construction is simpler than that of [DS11a] in that we do not need to concern ourselves with topological spanning trees (whose existence is nontrivial). On the other hand, our construction uses the *universal cover* of Γ_{hyp} , which is an internal tree where homotopies can be reparameterized in nice ways. Only with the hyperfinite approach is the recourse to the universal covering tree possible. Some applications of this result (and the construction of Θ) will be given in Section 5, where we give another proof of the fact that any inclusion of locally finite connected graphs $\Gamma^1 \rightarrow \Gamma^2$ induces an injection $\pi_1(|\Gamma^1|) \rightarrow \pi_1(|\Gamma^2|)$, the main reason being that the analogous result is true for finite graphs.

The original motivation for the work in [DS11a] was to show that certain loops in $|\Gamma|$ are non-nullhomologous. The motivation for this is that these

loops are trivial in the so-called topological cycle space, another homology theory that has proven to be more suitable than the usual one to study Fredudenthal compactifications (of spaces even more general than graphs) [DS10, DS11b]. The existence of such loops implies that these theories are actually different. Using our techniques and the algorithm for calculating commutator lengths in finitely generated free groups, in Section 5 we give a very short proof that one of the loops in the infinite-sideways ladder above witnessing the difference between the said homology theories is in fact non-nullhomologous. We also make the remark that a loop α gives rise to the trivial element of the topological cycle space if and only if $\Theta(\alpha)$, for Θ as in the main theorem, is internally null-homologous in $H_1(\Gamma_{\text{hyp}})$ (with coefficients in $\mathbb{Z}/2\mathbb{Z}$ or \mathbb{Z} depending on the version of topological cycle space under consideration).

We want to caution that our result does not imply that $\pi_1(|\Gamma|)$ is a free group. (According to Corollary 18 of [DS11a], $\pi_1(|\Gamma|)$ is free if and only if every end of Γ is contractible in $|\Gamma|$.) In general, *an internally free group is not actually free*. For example, \mathbb{Z}^* is internally free on one generator. Fix $M, N \in \mathbb{N}^* \setminus \mathbb{N}$ with $\frac{M}{N}$ infinitesimal and consider the injective group morphism $(a, b) \mapsto aM + bN : \mathbb{Z}^2 \rightarrow \mathbb{Z}^*$. Since \mathbb{Z}^2 is not free, we see that \mathbb{Z}^* is not free.

We would like to thank Reinhard Diestel for suggesting this project to us and for being so supportive of the nonstandard approach.

2. PRELIMINARIES ON END COMPACTIFICATIONS

In this section, we fix a proper, pointed geodesic metric space (X, d, p) . For $n \in \mathbb{N}$, we let $B(p, n)$ (resp. $\mathring{B}(p, n)$) denote the closed (resp. open) ball of radius n around p . For $x, y \in X$, we write $x \propto_n y$ to mean that there is a path in X connecting x to y which remains entirely in $X \setminus B(p, n)$. Given proper rays $r_1, r_2 : [0, \infty) \rightarrow X$ based at p , we say that r_1 and r_2 *determine the same end of X* if for every $n \in \mathbb{N}$, there is $t_0 \in [0, \infty)$ such that for all $t \geq t_0$, $r_1(t) \propto_n r_2(t)$. The notion of determining the same end is an equivalence relation on the set of proper rays beginning at p and we let $\text{end}(r)$ denote the equivalence class of the proper ray r .

The *end compactification of X* , denoted $|X|$, is the set X together with its ends. There is a natural topology on $|X|$, where a neighborhood basis of $x \in X$ is a neighborhood basis of x induced by the metric and a neighborhood basis for $\text{end}(r)$ is given by sets of the form

$$W_n(\text{end}(r)) = \{\text{end}(r') \mid \exists m_0 \in \mathbb{N} \forall m \geq m_0 r(m) \propto_n r'(m)\} \cup \\ \{x \in X \mid \exists m_0 \in \mathbb{N} \forall m \geq m_0 r(m) \propto_n x\}.$$

In [Gol11], a nonstandard characterization of the topology of the space of ends of X was outlined. In this section, we adapt the arguments from [Gol11] to show that $|X|$ is the same as the space $|X|_{\text{ns}}$ (the “nonstandard

ends compactification”) defined below. First, set

$$X_{\text{fin}} := \{x \in X^* \mid d(x, p) \text{ is finite}\} \quad \text{and} \quad X_{\text{inf}} := X^* \setminus X_{\text{fin}}.$$

For $x, y \in X^*$, write $x \propto y$ (resp. $x \propto_n y$ for some $n \in \mathbb{N}$) if there exists an internal path connecting x to y in $X^* \setminus X_{\text{fin}}$ (resp. in $X^* \setminus B(p, n)$). For $x \in X_{\text{inf}}$, let $[x]$ denote the equivalence class of x with respect to \propto . Set $\text{IPC}(X) := \{[x] \mid x \in X_{\text{inf}}\}$ and $|X|_{\text{ns}} := X \cup \text{IPC}(X)$. We consider the topology on $|X|_{\text{ns}}$ where a neighborhood basis for $[x] \in \text{IPC}(X)$ is given by

$$V_n([x]) := \{[x'] \in \text{IPC}(X) \mid x' \propto_n x\} \cup \{x' \in X \mid x' \propto_n x\},$$

and a neighborhood basis for $x \in X$ is a neighborhood basis for the topology induced by the metric.

In [Gol11], it is proven that $\text{IPC}(X)$ is homeomorphic to the space $\text{Ends}(X)$ of the ends of X via the homeomorphism $\eta : \text{Ends}(X) \rightarrow \text{IPC}(X)$ given by $\eta(\text{end}(r)) = [r(\sigma)]$, where $\sigma \in \mathbb{R}_{\text{inf}}^{>0}$ is fixed. Extend η to a map $\bar{\eta} : |X| \rightarrow |X|_{\text{ns}}$ by requiring $\bar{\eta}$ to be the identity map on X .

Proposition 2.1. $\bar{\eta} : |X| \rightarrow |X|_{\text{ns}}$ is a homeomorphism.

Proof. We begin with the following

Claim: $W_n(\text{end}(r)) = W'_n(\text{end}(r))$, where

$$W'_n(\text{end}(r)) = \{\text{end}(r') \mid r(\sigma) \propto_n r'(\sigma)\} \cup \{x \in X \mid x \propto_n r(\sigma)\}.$$

Proof of Claim: The containment \subseteq follows by transfer as $\sigma \geq m_0$. For the other containment, suppose that $\text{end}(r') \in W'_n(\text{end}(r))$, so $r(\sigma) \propto_n r'(\sigma)$. If $\text{end}(r') \notin W_n(\text{end}(r))$, then, by overflow, there is $\sigma' > \mathbb{N}$ such that $r(\sigma') \not\propto_n r'(\sigma')$. Since $r(\sigma) \propto_n r(\sigma')$ and $r'(\sigma) \propto_n r'(\sigma')$, we get $r(\sigma) \not\propto_n r'(\sigma)$, a contradiction. The other case is similar.

Since $|X|$ is compact and $\bar{\eta}$ is bijective, it suffices to show that $\bar{\eta}$ is continuous. $\bar{\eta}$ is clearly continuous on X , so let us consider $\text{end}(r) \in |X|$ and $n > 0$. The proof will be finished if we can show that $\bar{\eta}(W'_n(\text{end}(r))) = V_n([r(\sigma)])$. However, this is immediate from the definitions and the fact that η is a homeomorphism. \square

Lemma 2.2. If $|X^*|$ denotes the internal end compactification of X^* , then $|X|^* = |X^*|$.

Proof. Both spaces are the union of X^* and the quotient of the internally proper rays in X^* emanating from p modulo the relation of determining the same internal end. \square

3. EMBEDDING $\pi_1(\Gamma)$ INTO $\pi_1(\Gamma_{\text{hyp}})$

3.1. The Main Theorem. In this section, we fix an infinite, locally finite, connected graph Γ with end compactification $|\Gamma|$. We also fix a basepoint $p \in \Gamma$.

Lemma 3.1. *Suppose that $n \in \mathbb{N}$ and $\theta : \Gamma \rightarrow \Gamma_n$ is the map which collapses each connected component of $\Gamma \setminus \mathring{B}(p, n)$ to a point. Then θ extends continuously to a map $\theta : |\Gamma| \rightarrow \Gamma_n$.*

Proof. Given $\text{end}(r) \in |\Gamma|$ and $x \in \Gamma \cap W_n(\text{end}(r))$, set $\theta(\text{end}(r)) := \theta(x)$. It is then easy to check that $\theta(W_n(r)) = \{\theta(x)\}$, whence θ is continuous at $\text{end}(r)$. \square

In the rest of this section, we fix $\nu \in \mathbb{N}^* \setminus \mathbb{N}$ and set $\theta : \Gamma^* \rightarrow \Gamma_{\text{hyp}}$ to be the map which collapses each internally connected component of $\Gamma^* \setminus \mathring{B}(p, \nu)$ to a point. We consider Γ both as a subset of $|\Gamma|^*$ and Γ_{hyp} in the obvious way and let $\hat{\Gamma}$ denote both the elements of $|\Gamma|^*$ and Γ_{hyp} an infinitesimal distance away from an element of Γ . (This double use of $\hat{\Gamma}$ should not cause any confusion.) By the previous lemma, we obtain an extension of θ to an internally continuous map $\theta : |\Gamma^*| = |\Gamma|^* \rightarrow \Gamma_{\text{hyp}}$. Consequently, we get an induced map on the internal fundamental groups:

Definition 3.2. We set $\Theta : \pi_1(|\Gamma|^*) \rightarrow \pi_1(\Gamma_{\text{hyp}})$ to be the group homomorphism induced by the map $\theta : |\Gamma|^* \rightarrow \Gamma_{\text{hyp}}$. Here, $\pi_1(|\Gamma|^*)$ and $\pi_1(\Gamma_{\text{hyp}})$ are the internal fundamental groups of $|\Gamma|^*$ and Γ_{hyp} respectively.

Observe that $\pi_1(|\Gamma|^*) = (\pi_1(|\Gamma|))^*$, so $\pi_1(|\Gamma|)$ is a subgroup of $\pi_1(|\Gamma|^*)$. Our main goal is to show that $\Theta|_{\pi_1(|\Gamma|)}$ is injective, proving Theorem 1.1. First, we show that we can construct $|\Gamma|$ in a nonstandard fashion analogous to the construction of Section 2 using Γ_{hyp} instead of Γ^* .

Here, we are slightly abusive in notation and let $\hat{\Gamma}$ refer to the elements of either Γ^* or Γ_{hyp} that are infinitely close to an element of Γ .

We let d denote the usual metric on Γ , namely the path metric on Γ when each edge is identified with the interval $[0, 1]$. Similarly, we let d_h denote the internal metric on Γ_{hyp} obtained by identifying each edge with $[0, 1]^*$.

For $x, y \in \Gamma_{\text{hyp}}$, write $x \propto y$ (resp. $x \propto_n y$ for some $n \in \mathbb{N}$) if there exists an internal path connecting x to y in $\Gamma_{\text{hyp}} \setminus \hat{\Gamma}$ (resp. in $\Gamma_{\text{hyp}} \setminus B(p, n)$), and denote by $[x]$ the equivalence class of x with respect to \propto . Also, let $\text{IPC}_{\text{hyp}}(\Gamma)$ denote the collection of such equivalence classes. Finally, consider the topology on $|\Gamma|_{\text{hyp}} := \Gamma \cup \text{IPC}_{\text{hyp}}(\Gamma)$ where a neighborhood basis for $[x] \in \text{IPC}_{\text{hyp}}(\Gamma)$ is given by

$$V_n^{\text{hyp}}([x]) := \{[x'] \in \text{IPC}_{\text{hyp}}(\Gamma) \mid x' \propto_n x\} \cup \{x' \in \Gamma \mid x' \propto_n x\},$$

and a neighborhood basis for $x \in \Gamma$ is a neighborhood basis for the topology induced by the metric.

Suppose $x, y \in \Gamma_{\text{inf}}$ are such that there is an internal path connecting x to y in $\Gamma^* \setminus \hat{\Gamma}$. Then, composing with θ , we get an internal path connecting $\theta(x)$ to $\theta(y)$ in $\Gamma_{\text{hyp}} \setminus \hat{\Gamma}$. This allows us to define a map $\chi : |\Gamma|_{\text{ns}} \rightarrow |\Gamma|_{\text{hyp}}$ which restricts to the identity on Γ and satisfies $\chi([x]) = [\theta(x)]$ for each $x \in \Gamma_{\text{inf}}$.

Lemma 3.3. χ is a homeomorphism.

Proof. Arguing as before the lemma, $V_n([x])$ is mapped into $V_n^{\text{hyp}}([\theta(x)])$ for each $x \in \Gamma_{\text{inf}}$, whence χ is continuous. Since $|\Gamma|_{\text{ns}}$ is compact, it suffices to prove that χ is bijective. Since χ is clearly surjective, it remains to prove that χ is injective. Towards this end, suppose that $\theta(x)$ can be connected to $\theta(y)$ by an internal path γ in $\Gamma_{\text{hyp}} \setminus \hat{\Gamma}$. Consider pairs of values t_0^j, t_1^j so that $d_h(\gamma(t_0^j), p) = \nu$, $\gamma(t_0^j)$ is in the same connected component of $\Gamma^* \setminus \mathring{B}(p, \nu)$ as $\gamma(t_1^j)$, and the interval $[t_0^j, t_1^j]$ is maximal among intervals with endpoints satisfying these properties. We can substitute $\gamma|_{[t_0^j, t_1^j]}$ with an internal path in $\Gamma^* \setminus \mathring{B}(p, \nu)$ connecting $\gamma(t_0^j)$ to $\gamma(t_1^j)$ (if $t_0^j > 0$ and $t_1^j < 1$, otherwise use x and/or y), and so we obtain an internal path from x to y in $\Gamma^* \setminus \hat{\Gamma}$. \square

Theorem 3.4. $\Theta|_{\pi_1(|\Gamma|)} : \pi_1(|\Gamma|) \rightarrow \pi_1(\Gamma_{\text{hyp}})$ is injective.

The following corollary will be useful for our applications.

Corollary 3.5. *If Λ is a hyperfinite graph and there exist internally continuous maps $\varphi : |\Gamma|^* \rightarrow \Lambda$, $\psi : \Lambda \rightarrow \Gamma_{\text{hyp}}$ so that $\psi \circ \varphi = \theta$, then the map induced by φ at the level of fundamental groups is injective when restricted to $\pi_1(|\Gamma|)$.*

Proof. (of Theorem 3.4) Suppose that $\theta(\alpha)$ is (internally) null-homotopic, for some loop $\alpha : [0, 1] \rightarrow |\Gamma|$ based at $p \in \Gamma$. We will construct a continuous family of paths γ_t connecting p to $\alpha(t)$. The idea for constructing such a family is the following. As $\theta(\alpha)$ is null-homotopic, it lifts to an internal loop in the internal universal cover $\widetilde{\Gamma}_{\text{hyp}}$ of Γ_{hyp} , which is an internal tree. We can then project on Γ_{hyp} a family of geodesics connecting a lift of p to points in the lift of $\theta(\alpha)$ based at the same point, and use these projected paths to construct the required homotopy.

Convenient reparameterizations. Even though it is not absolutely necessary, we will assume that α makes no partial crossing of edges, meaning that if the point $\alpha(x)$ is contained in the interior of an edge then there is an interval I containing x so that $\alpha(I)$ is the edge containing $\alpha(x)$ and $\alpha|_I$ connects the endpoints of the said edge. We can achieve this by a homotopy in view of [DS11a, Lemma 2]. Let $\rho : \mathbb{N}^* \rightarrow \mathbb{N}^*$ be the function such that $\rho(\xi)$ is the cardinality of the edges at distance at most ξ from $p \in \Gamma_{\text{hyp}}$ that are crossed by $\theta(\alpha)$, counted with multiplicity. (The distance of an edge from p is the minimal distance of its endpoints from p .) Notice that if ξ is finite, then $\rho(\xi)$ is finite as it coincides with the number of edges at distance at most ξ from p crossed by α counted with multiplicity. We can reparameterize $\theta(\alpha) : [0, \lambda] \rightarrow \Gamma_{\text{hyp}}$, for some finite $\lambda \in \mathbb{R}^*$, so that subintervals of $[0, \lambda]$ which correspond to crossing an edge at distance ξ from p have length $1/(2^\xi \rho(\xi))$. Notice that λ is indeed finite because, as there are at most $\rho(\xi)$ subpaths of $\theta(\alpha)$ corresponding to crossing an edge at distance ξ from p , we have $\lambda \leq \sum_{\xi \in \mathbb{N}^*} \rho(\xi)/(2^\xi \rho(\xi)) = 1$. Also, we can reparameterize α so that for each $\tau \in [0, \lambda]$ we have $\alpha(\text{st}(\tau)) = \text{st}(\theta(\alpha)(\tau))$ if $\theta(\alpha)(\tau) \in \hat{\Gamma}$, and $\alpha(\text{st}(\tau))$ is the end corresponding to $\theta(\alpha)(\tau)$ otherwise

(the correspondence is given by Lemma 3.3). Here, $\text{st} : \hat{\Gamma} \rightarrow \Gamma$ is the map so that $\text{st}(x)$ has infinitesimal distance from x .

The homotopy. Let $f : \widetilde{\Gamma}_{\text{hyp}} \rightarrow \Gamma_{\text{hyp}}$ be the universal covering map and fix $\tilde{p} \in \widetilde{\Gamma}_{\text{hyp}}$ such that $f(\tilde{p}) = p$. As mentioned above, the fact that $\theta(\alpha)$ is null-homotopic means that it can be lifted to a loop $\widetilde{\theta(\alpha)}$ in $\widetilde{\Gamma}_{\text{hyp}}$ based at \tilde{p} , that is, $f \circ \widetilde{\theta(\alpha)} = \theta(\alpha)$. Let \tilde{d} the metric of $\widetilde{\Gamma}_{\text{hyp}}$. Fix $t \in [0, \text{st}(\lambda)]$, let $\tilde{\delta}_t : [0, \tilde{d}(\tilde{p}, \widetilde{\theta(\alpha)}(t))] \rightarrow \widetilde{\Gamma}_{\text{hyp}}$ be the geodesic connecting \tilde{p} to $\widetilde{\theta(\alpha)}(t)$ and let $\delta_t := f \circ \tilde{\delta}_t$. (It is possible that $\lambda < \text{st}(\lambda)$, whence these definitions are ill-defined for $t = \text{st}(\lambda)$; in this case, let $\tilde{\delta}_{\text{st}(\lambda)}$ be the geodesic connecting \tilde{p} to $\widetilde{\theta(\alpha)}(\lambda)$ and let $\delta_{\text{st}(\lambda)} := f \circ \tilde{\delta}_{\text{st}(\lambda)}$.) Reparameterize δ_t to obtain $\phi_t : [0, \lambda_t] \rightarrow \Gamma_{\text{hyp}}$ with the property that subintervals of $[0, \lambda_t]$ which correspond to crossing an edge at distance ξ from p have length $1/(2^\xi \rho(\xi))$. Notice that the image of δ_t (and so that of ϕ_t) is contained in that of $\theta(\alpha)$. In fact, given any path in a tree, the (image of the) geodesic connecting its endpoints is contained in (the image of) the path. As a consequence, ϕ_t is the concatenation of subpaths of $\theta(\alpha)$ with disjoint domain and so λ_t is finite.

Define $\gamma_t : [0, 1] \rightarrow |\Gamma|$ by $\gamma_t(u) = \text{st}(\phi_t(u\lambda_t))$ if $\phi_t(u\lambda_t) \in \hat{\Gamma}$, and let $\gamma_t(u)$ be the end corresponding to $\phi_t(u\lambda_t)$ otherwise (we are using Lemma 3.3 again).

We claim that $H : [0, \text{st}(\lambda)] \times [0, 1] \rightarrow |\Gamma|$ defined by $H(t, u) := \gamma_t(u)$ is continuous and hence provides a homotopy from α to the trivial loop.

Long common subpaths. The key observation is that ϕ_{t_0}, ϕ_{t_1} differ at most in final intervals of the respective domains of length bounded by $|t_0 - t_1|$, meaning that we have $\phi_{t_0}|_{[0, \lambda]} = \phi_{t_1}|_{[0, \lambda]}$ for $\lambda = \max\{\lambda_{t_i} - (t_1 - t_0)\}$. To see this, suppose, without loss of generality, that $t_0 < t_1$ and let β be the maximal common initial subpath of ϕ_{t_0} and ϕ_{t_1} ; we must show that the length of the domain of β is at least $\lambda_{t_i} - (t_1 - t_0)$ for $i = 0, 1$. For $i = 0, 1$, let ϕ'_i denote the final subpath of ϕ_{t_i} connecting the last point of β to $\theta(\alpha)(t_i)$. Then the lift of ϕ'_i is contained in the geodesic segment connecting $\widetilde{\theta(\alpha)}(t_0)$ to $\widetilde{\theta(\alpha)}(t_1)$, whose image in turn is contained in the image of $\theta(\alpha)|_{[t_0, t_1]}$. As explained above, this means that the image of ϕ'_i is contained in the image of $\theta(\alpha)|_{[t_0, t_1]}$ and ϕ'_i is a concatenation of subpaths of $\theta(\alpha)|_{[t_0, t_1]}$, whence the domain of ϕ'_i is at most $t_1 - t_0$. Thus the claim is proven.

Proof of continuity. Fix $t \in [0, \text{st}(\lambda)]$ and $u \in [0, \text{st}(\lambda_t)]$. Define $\gamma'_t(u) = \text{st}(\phi_t(u))$ if $\phi_t(u) \in \hat{\Gamma}$ and $\gamma'_t(u)$ is the end corresponding to $\phi_t(u)$ otherwise. (Once again, this formula may be ill-defined for $u = \text{st}(\lambda_t)$ if $\lambda_t < \text{st}(\lambda_t)$; in this case, we make the obvious change to the definition of $\gamma'_t(\text{st}(\lambda_t))$.) We now consider the function

$$H' : \{(t, u) \in [0, \text{st}(\lambda)] \times \mathbb{R}_{\geq 0} : u \leq \text{st}(\lambda_t)\} \rightarrow |\Gamma|$$

given by $H'(t, u) := \gamma'_t(u)$. Since the key observation implies that $t \mapsto \text{st}(\lambda_t)$ is continuous, in order to check continuity of H it suffices to check continuity of H' . (Notice that $\phi_t(u \cdot \text{st}(\lambda_\tau))$ is infinitely close to $\phi_t(u \cdot \lambda_t)$ if one of them is in $\hat{\Gamma}$ and that otherwise the subpath of ϕ_t connecting them does not intersect $\hat{\Gamma}$ so that they correspond to the same end.)

We first show that H' is continuous at (t, u) if $H'(t, u) \in \Gamma$. Indeed, we claim that for $n \geq d(p, \phi_t(u)) + 1$ and $\epsilon < 1/2$, if $|t - t'| < \epsilon/(2^{n+1}\rho(n+1))$ and $|u - u'| < \epsilon/(2^{n+1}\rho(n+1))$, then we have $d(H'(t, u), H'(t', u')) \leq 3\epsilon$. To see this, first observe that if $|t - t'| < \epsilon/(2^{n+1}\rho(n+1))$, then ϕ_t and $\phi_{t'}$ either differ by paths contained in $B(p, n+1)$ whose lengths are each bounded above by ϵ or by paths contained in the complement of $B(p, n-1)$. Indeed, taking into account our choice of parameterization, if the endpoint of ϕ_t is in $B(p, n)$, then $\phi_t|_{[\max\{0, \lambda_t - |t - t'|, \lambda_t\}]}$ is contained in $B(p, n+1)$ and hence has length at most ϵ , and otherwise the said subpath cannot travel through an edge at distance $n-1$ from p . To finish the proof of our claim, it suffices to observe that if β is any path contained in $B(p, n+1)$ satisfying our choice of the parameterizations (e.g. $\phi_t, \phi_{t'}$) and if $|u - u'| < \epsilon/(2^{n+1}\rho(n+1))$, then $d_h(\beta(u), \beta(u')) < \epsilon$.

Suppose instead that $H'(t, u)$ is an end. We claim that if $|t - t'| < 1/(2^{n+1}\rho(n+1))$ and $|u - u'| < 1/(2^{n+1}\rho(n+1))$, then $\phi_t(u)$ can be connected to $\phi_{t'}(u')$ by a path lying outside $B(p, n)$, which proves the continuity of H' at (t, u) in view of the nonstandard characterization of the topology of $|\Gamma|$ and the fact that $\phi_t(u), \phi_{t'}(u')$ represent the same point of $|\Gamma|$ as $H'(t, u), H'(t', u')$.

With the given restrictions on t', u' , we have that $\phi_t, \phi_{t'}$ differ by paths contained in some connected component C_1 of $\Gamma_{\text{hyp}} \setminus B(p, n)$ and also $\phi_{t'}(u), \phi_{t'}(u')$ must lie in the same connected component C_2 of $\Gamma_{\text{hyp}} \setminus B(p, n)$. If $C_1 \neq C_2$ then ϕ_t and $\phi_{t'}$ coincide in C_2 , so the path connecting $\phi_t(u)$ to $\phi_{t'}(u')$ outside $B(p, n)$ can be taken to be a subpath of ϕ_t . If $C_1 = C_2$ then by definition $\phi_{t'}(u'), \phi_{t'}(u)$ and $\phi_{t'}(u), \phi_t(u)$ lie in the same connected component of $\Gamma_{\text{hyp}} \setminus B(p, n)$. \square

Remark 3.6. The combinatorial characterization of $\pi_1(|\Gamma|)$ given in [DS11a] also applies to certain subspaces of $|\Gamma|$, the so-called *standard subspaces* of $|\Gamma|$. The subspace H of $|\Gamma|$ is called *standard* if H is closed, connected, and contains every edge of which it contains an inner point. We leave it to the reader to check that our construction is readily adaptable to the case of a standard subspace of $|\Gamma|$.

3.2. The image of Θ . It would be nice if we could characterize the image of our embedding $\Theta|_{\pi_1(|\Gamma|)}$, but unfortunately we are currently unable to accomplish this goal. In this subsection, we discuss some of the difficulties we face in this endeavor.

Once again, let us consider the 1-way sideways infinite ladder Γ from the Introduction. Below, in Figure 3, we have Γ_{hyp} with its internal spanning tree and chords e_1, e_2, \dots, e_ν .




FIGURE 3

The question we face is which internal words are in the image of the embedding $\Theta : \pi_1(|\Gamma|) \rightarrow \pi_1(\Gamma_{\text{hyp}})$? Here is a word that is in the image of this embedding: $e_1 e_2 \cdots e_\nu$. Indeed, this word arises from the loop α that travels down the bottom edge of the ladder to the unique end and then returns along the top edge of the ladder. However, here is a word that is *not* in the image of the embedding: $e_1 e_2 \cdots e_{\nu-1}$. Indeed, in some sense, this word should result from the same loop that gave rise to $e_1 e_2 \cdots e_\nu$, but it is impossible for the standard loop to “stop” at $e_{\nu-1}$ and continue upwards. More generally, for any $\eta \in \mathbb{N}^* \setminus \mathbb{N}$ with $\eta < \nu$, the word $e_1 e_2 \cdots e_\eta$ should not be in the image of the embedding.

We also need to ensure that words representing “loops at infinity” are not in the image of our embedding. For example, consider the word

$$e_1 e_2 \cdots e_{\nu-1} e_\nu e_{\nu-1}^{-1} \cdots e_2^{-1} e_1^{-1}.$$

Indeed, if this word arose from some loop α , then by transfer, α would have made a loop at some finite stage $n \in \mathbb{N}$ and the image of $[\alpha]$ under our embedding would have been $e_1 e_2 \cdots e_{n-1} e_n e_{n-1}^{-1} \cdots e_2^{-1} e_1^{-1}$. With these hurdles in mind, we pose the following

Question 3.7. *What is the image of the embedding $\Theta : \pi_1(|\Gamma|) \rightarrow \pi_1(\Gamma_{\text{hyp}})$?*

4. CONNECTION TO THE RESULTS IN [DS11a]

4.1. Embedding into an inverse limit of free groups. In this subsection, we observe that $\Theta : \pi_1(|\Gamma|) \rightarrow \pi_1(\Gamma_{\text{hyp}})$ factors through an injective map $\Psi : \pi_1(|\Gamma|) \rightarrow \varprojlim F_n$ for certain finitely generated free groups F_n . For $n \in \mathbb{N}$, set $\theta_n : \Gamma \rightarrow \Gamma_n$ to be the map collapsing each connected component of $\Gamma \setminus \mathring{B}(p, n)$ to a point. Set $F_n := \pi_1(\Gamma_n)$. There are natural continuous maps $\rho_n^m : \Gamma_m \rightarrow \Gamma_n$ when $m \geq n$, so that $\{F_n\}$ is an inverse system of groups. By Lemma 3.1, we have a continuous extension of θ_n to a map $\theta_n : |\Gamma| \rightarrow \Gamma_n$ satisfying $\theta_n = \rho_n^m \circ \theta_m$ whenever $m \geq n$. We thus have a map $\Psi : \pi_1(|\Gamma|) \rightarrow \varprojlim F_n$.

Suppose we choose our nonstandard extension Γ^* to be $\Gamma^{\mathbb{N}}/\mu$ for some non-principal ultrafilter μ on \mathbb{N} , and choose $\nu = [(\text{id} : \mathbb{N} \rightarrow \mathbb{N})]_\mu$. Then $\Gamma_{\text{hyp}} = \prod_\mu \Gamma_n$. Consequently, $\pi_1(\Gamma_{\text{hyp}}) = \prod_\mu F_n$. Now consider the map $\Phi : \varprojlim F_n \rightarrow \prod_\mu F_n$ given by $\Phi((x_n)) = [(x_n)]_\mu$. Then Φ is injective: if $x_n = y_n$ μ -a.e., then since μ is nonprincipal, $x_n = y_n$ for infinitely many n . Since $(x_n), (y_n) \in \varprojlim F_n$, it follows that $x_n = y_n$ for all n .

We leave it to the reader to check that $\Theta = \Phi \circ \Psi$. It follows that Ψ is injective.

4.2. The Standard Description of $\pi_1(|\Gamma|)$. After seeing an initial draft of this article, Diestel asked us whether we could recover the main theorem of [DS11a], namely Theorem 15. After some further effort, we can indeed accomplish this. First, let us briefly describe the content of Theorem 15. Fix a topological spanning tree T of $|\Gamma|$, namely an arc-connected, closed subspace of $|\Gamma|$ that does not contain any circle (that is, homeomorphic image of \mathbb{S}^1) and which contains every edge of which it contains an interior point. We let $\{e_i : i \in \mathbb{N}\}$ enumerate the chords of T , namely the edges of Γ not in T , equipped with a fixed orientation. (We may assume that there are infinitely many chords, for otherwise $|\Gamma|$ is homotopy equivalent to a finite graph and there is nothing to prove.)

Given a loop α in $|\Gamma|$ based at p , one can consider the trace of α on the chords, giving a countable “word” w_α , which should be viewed as a function from some countable linearly ordered set S to the set $A := \{e_i^{\pm 1} : i \in \mathbb{N}\}$. The authors of [DS11a] define a suitable notion of reduction of words and prove that every word w reduces to a unique reduced word $r(w)$. They let F_∞ denote the set of reduced words and endow it with the obvious group multiplication (namely $r(w) \cdot r(w')$ is the unique reduced word associated to the concatenation of w and w'). They prove that if α and β are homotopic loops based at p , then $r(w_\alpha) = r(w_\beta)$, whence they obtain a map $[\alpha] \mapsto r(w_\alpha) : \pi_1(|\Gamma|) \rightarrow F_\infty$. Furthermore, they prove that this map is *injective* and characterize the image of this map as precisely the set of words whose monotonic subwords converge (we will explain this terminology later in the next subsection). They then show that F_∞ embeds into an inverse limit of finitely generated free groups by mapping $r(w)$ to the family $(r(w \upharpoonright I))_{I \subseteq S \text{ finite}}$. They characterize the image of this latter embedding as the set of elements of the inverse limit such that each letter appears only finitely often. As we remarked in the introduction, this latter embedding is purely algebraic and the nonstandard set-up has nothing to say about this portion of the theorem.

By far the hardest part of the paper [DS11a] is their Lemma 14, which shows that if $r(w_\alpha) = \emptyset$ for a loop α based at p , then α is nullhomotopic, implying the injectivity of the map $\pi_1(|\Gamma|) \rightarrow F_\infty$. In what follows, we use our framework to give another proof of Lemma 14. Our proof is significantly shorter than the standard proof, which is 14 pages in length.

4.3. Recovering The Standard Description. We let Λ' be the (possibly non-simplicial) internal graph obtained collapsing the internal connected components of $T^* \setminus \mathring{B}(p, \nu)$ to points. Also, let Λ be the (simplicial) graph obtained collapsing loops in Λ' to points and identifying pairs of edges with the same endpoints. We let $\varphi : |\Gamma|^* \rightarrow \Lambda$ be the natural internally continuous map. Also, we let $T' = \varphi(T^*)$, and we remark that T' is an internal tree.

Notice that if e is a chord of T^* , then either φ maps e to a vertex in Λ or else it maps e to a chord of T' .

Lemma 4.1. *Given any chord e of T^* so that $\varphi(e)$ is a chord of T' , there are only hyperfinitely many chords e' of T^* with $\varphi(e) = \varphi(e')$.*

Proof. If not, there are two connected components T_1, T_2 of $T^* \setminus \mathring{B}(p, \nu)$ that are connected by more than hyperfinitely many edges. By transfer, there would then exist an integer n and two connected components T_1, T_2 of $T \setminus \mathring{B}(p, n)$ connected by infinitely many edges. But then we would have that T_1 and T_2 have a common end in both of their closures, so they were not different connected components after all. \square

Let $\theta : |\Gamma|^* \rightarrow \Gamma_{\text{hyp}}$ be as in Section 3. Observe that θ factors as $\theta = \psi \circ \varphi$, where $\psi : \Lambda \rightarrow \Gamma_{\text{hyp}}$ maps the point of Λ corresponding to the internal connected component C of $T^* \setminus \mathring{B}(p, \nu)$ to the internal connected component of $|\Gamma|^* \setminus \mathring{B}(p, \nu)$ containing C . Since θ is π_1 -injective, it follows that φ is π_1 -injective (Corollary 3.5).

Proposition 4.2 ([DS11a], Lemma 14). *Suppose that $\alpha : [0, 1] \rightarrow |\Gamma|$ is such that w_α reduces to the empty word. Then α is nullhomotopic.*

Proof. Set $w := w_\alpha$ and say $w : S \rightarrow A$, where S is some countable, linearly ordered set. Let $S_0 \subseteq S$ be the set of all $s \in S$ so that $\phi(w(s))$ is a chord of T' . Notice that S is hyperfinite because of the hyperfiniteness of Λ , Lemma 4.1 and the fact that for each $e \in A$ there can be only hyperfinitely many $s \in S$ so that $w(s) = e$. We have a reduction R of w to the empty word, whence R^* is an internal reduction of w^* to the empty word. By (1) on page 8 of [DS11a], we can restrict such reduction to an internal reduction of $w^* \upharpoonright S_0$ to the empty word. Since $w^* \upharpoonright S_0$ is a hyperfinite word, $w^* \upharpoonright S_0$ internally reduces (in the usual sense) to the empty word. But then, considering the trace of $\varphi(\alpha)$ on the chords of T' , we see that $\varphi(\alpha)$ is internally nullhomotopic. Therefore, α is nullhomotopic. \square

We now proceed to prove the other difficult part of Theorem 15, namely the part which characterizes the image of the embedding $\pi_1(|\Gamma|) \rightarrow F_\infty$. To explain this, we need some terminology. First, given a word $w : S \rightarrow A$, we say that a subword $w \upharpoonright S' : S' \rightarrow A$ is *monotonic* if S' is infinite and we can write $S' := \{s_0, s_1, s_2, \dots\}$ where either $s_0 < s_1 < s_2 < \dots$ or $s_0 > s_1 > s_2 > \dots$. We say that the monotonic subword $w \upharpoonright S'$ of w *converges* if there exists an end $\text{end}(r)$ of $|\Gamma|$ such that, whenever we choose $x_n \in w(s_n)$, the sequence (x_n) converges to $\text{end}(r)$. It is very easy to see that, for any loop α in $|\Gamma|$ based at p , every monotonic subword of the word w_α converges.

Proposition 4.3 (Part of [DS11a], Theorem 15). *Suppose that the word $w : S \rightarrow A$ has the property that all its monotonic subwords converge. Then there exists a loop α so that $w = w_\alpha$.*

Proof. We fix Λ and T' as above. As in the proof of the previous proposition, we consider the hyperfinite word $w^* \upharpoonright S_0$, where $S_0 \subseteq S$ is the set of all $s \in S$ so that $\phi(w(s))$ is a chord of T' . As in the case of finite graphs, $w^* \upharpoonright S_0$ corresponds to an internal loop β in Λ based at p obtained by alternately concatenating chords of T' (given by $\varphi(e)$ for e running through the letters of $w^* \upharpoonright S_0$) and injective paths in T' .

Claim: β crosses each edge of $\Gamma \subseteq \Lambda$ finitely many times.

For the moment, suppose that the Claim holds. As we have done several times, we consider $\rho : \mathbb{N}^* \rightarrow \mathbb{N}^*$, where $\rho(\xi)$ is the the number of edges at distance at most ξ from the basepoint p that β crosses, counted with multiplicity. The claim implies that $\rho(n) \in \mathbb{N}$ for $n \in \mathbb{N}$. We then reparameterize β in such a way that intervals in its domain corresponding to crossing an edge at distance ξ have length $2^{-\xi}\rho(\xi)$ (and such crossings are performed linearly). Consider the path α defined by $\alpha(t) = \text{st}(\beta(t))$ if $\text{st}(\beta(t))$ is defined, while $\alpha(t)$ is the end corresponding to $\psi(\beta(t))$ otherwise (we are using Lemma 3.3). Arguments identical to those earlier in the paper show that α is a loop in $|\Gamma|$ based at p and that $w = w_\alpha$, as we desired.

Proof of Claim: We must show that β crosses e finitely many times. This is easily seen to be the case if $e = \varphi(e_0)$ is a chord of T (notice that e_0 is unique). Indeed, we then have that e_0 and e_0^{-1} appear only finitely many times in w , for otherwise, by Ramsey's theorem, there would exist a monotonic subword $w \upharpoonright S'$ of w such that $w(s) \in \{e_0, e_0^{-1}\}$ for all $s \in S'$; it is clear that $w \upharpoonright S'$ does not converge. Consequently, e_0 and e_0^{-1} appear in $w^* \upharpoonright S_0$ only finitely many times, whence β crosses e only finitely many times.

We now suppose that $e \in T$ and further suppose, towards a contradiction, that e is crossed infinitely many times by β . First observe that $T' \setminus \dot{e}$ has two components, say T_1 and T_2 . For later purposes, notice that φ maps the connected components of $T^* \setminus \dot{e}$ to T_1 and T_2 . Given a chord of T' , we will say it is of type (i, j) , where $i, j \in \{1, 2\}$, if its initial point is in T_i and its final point is in T_j . Moreover, we say that a chord c of T^* is of type (i, j) if $\varphi(c)$ is a chord of T' of type (i, j) . We can then find a hyperfinite sequence s_0, s_1, \dots, s_ξ , where $s_i \in S^*$ and $\xi > \mathbb{N}$, so that going in T' from the final point of $\varphi(w(s_i))$ to the initial point of $\varphi(w(s_{i+1}))$ requires crossing e , whence $w(s_i)$ is of type $(*, k)$ and $w(s_{i+1})$ is of type (l, \dagger) , where $k \neq l$. Notice that going in T^* from the final point of $w(s_i)$ to the initial point of $w(s_{i+1})$ also requires crossing e . Indeed, if the unique path in T^* connecting the final point of $w(s_i)$ to the initial point of $w(s_{i+1})$ did not pass through e , then after applying φ , we would see that, since φ acts identically on Γ , the unique path in T' connecting the final point of $w(s_i)$ to the initial point of $w(s_{i+1})$ does not pass through e . It remains to note that the unique paths in T^* and T' connecting the final point of $w(s_i)$ to the initial point of $w(s_{i+1})$ agree on Γ .

We say that a subword of $w(s_0)w(s_1)\cdots w(s_\xi)$ is *uniform* if either every $w(s_i)$ has type $(1,2)$ or every $w(s_i)$ has type $(2,1)$, and we say that it is *alternating* if whenever $w(s_i)$ is of type $(1,1)$ (resp. $(2,2)$), then $w(s_{i+1})$ is of type $(2,2)$ (resp. $(1,1)$).

We claim that there is an infinite, hyperfinite subword of $w(s_0)w(s_1)\cdots w(s_\xi)$ that is either uniform or alternating. Indeed, either the subword consisting of chords of type $(1,2)$ is infinite or the subword consisting of chords of type $(2,1)$ is infinite, or else the subword obtained by removing all chords of type $(1,2)$ and $(2,1)$ is infinite and hyperfinite. We may thus suppose that $w(s_0)w(s_1)\cdots w(s_\xi)$ is either uniform or alternating.

Thus, by transfer, there exist arbitrarily long finite subwords $w \upharpoonright S'$ that are either uniform for e or alternating for e . If we have arbitrarily long finite uniform subwords, then we have an infinite uniform subword obtained by taking the union of the finite subwords. After applying Ramsey's theorem, we have an infinite, monotonic subword of w such that all letters correspond to an edge with an endpoint in T_1 and an endpoint in T_2 . Such a subword does not converge to an end as every end has a neighborhood that is disjoint from one of the components of $T \setminus \dot{e}$ (for otherwise T contains a circle).

Suppose instead that we have arbitrarily long finite alternating subwords. Let $S_1 \subseteq S$ be the set of all $s \in S$ so that $w(s)$ has type either $(1,1)$ or $(2,2)$, and let S'_1 be the set of all maximal intervals of S_1 that do not contain any s with $w(s)$ of type $(2,2)$. It is readily checked that S'_1 inherits a total order from S_1 . Also, under our assumption on alternating subwords, it is infinite. In particular, by Ramsey's theorem, there exists a monotonic subword $w \upharpoonright \{s'_0, \dots\}$ so that for each letter s'_i we have that $w(s'_i)$ has type $(1,1)$, and between s'_i and s'_{i+1} there exists some element of S of type $(2,2)$. By inserting such elements once again we find a monotonic subword of w such that there are infinitely many endpoints of chords in T_1 and infinitely many endpoints of chords in T_2 , and we can conclude as above. \square

5. APPLICATIONS

5.1. π_1 -injectivity. Both of the applications in this subsection rely on applying the transfer principle to the following

Fact 5.1. *If Γ^1 is a connected subgraph of the finite, connected graph Γ^2 , then the natural map $\pi_1(\Gamma^1) \rightarrow \pi_1(\Gamma^2)$ is injective.*

To see this, take a spanning tree T^1 of Γ^1 and extend it to a spanning tree T^2 of Γ^2 . Using these spanning trees in the combinatorial characterization of $\pi_1(\Gamma^1)$ and $\pi_1(\Gamma^2)$ yields the aforementioned result.

Proposition 5.2. *For any locally finite graph Γ , the inclusion $\Gamma' \rightarrow |\Gamma|$ is π_1 -injective for every subgraph Γ' of Γ .*

Proof. Let α be a non-trivial loop in Γ' . Then α is non-trivial in a finite connected subgraph Δ of Γ . Hence $\theta(\alpha) = \alpha$ is non-trivial in Δ^* and we can conclude because the inclusion $\Delta^* \rightarrow \Gamma_{\text{hyp}}$ is π_1 -injective by Fact 5.1. \square

Proposition 5.3 ([DS11a], Corollary 16). *Suppose that Γ^1 is a connected subgraph of the locally finite, connected graph Γ^2 . Then the inclusion $\iota : \Gamma^1 \rightarrow \Gamma^2$ induces an injective map $\iota_* : \pi_1(|\Gamma^1|) \rightarrow \pi_1(|\Gamma^2|)$.*

Proof. Let $\{\Delta_i\}$ be the collection of all internally connected components of $(\Gamma^1)^* \cap \left((\Gamma^2)^* \setminus \mathring{B}_{\Gamma^2}(p, \nu) \right)$. Let $\varphi : (\Gamma^1)^* \rightarrow \Lambda$ be the map that collapses each Δ_i to a point and let $\psi : \Lambda \rightarrow \Gamma_{\text{hyp}}^1$ be the map which further collapses the rest of the internal path components of $(\Gamma^1)^* \setminus B_{\Gamma^1}(p, \nu)$ to points. Since $\psi \circ \varphi = \theta_1$ (the usual collapsing map for $(\Gamma^1)^*$), we see that $\Phi|_{\pi_1(|\Gamma^1|)}$ is injective. We have a natural map $\iota_{\text{hyp}} : \Lambda \rightarrow \Gamma_{\text{hyp}}^2$ induced by ι , namely $\iota_{\text{hyp}}(\varphi(x)) = \theta_2(\iota(x))$. This map satisfies $(\iota_{\text{hyp}})_* \circ \Phi = \Theta_2 \circ \iota_*$ at the level of the fundamental groups, as can easily be checked. It thus suffices to show that ι_{hyp} is π_1 -injective. Towards this end, consider the graph Λ' obtained from Λ by identifying $x, y \in \Lambda$ if $\iota_{\text{hyp}}(x) = \iota_{\text{hyp}}(y)$. Notice that if $x \neq y$ and $\iota_{\text{hyp}}(x) = \iota_{\text{hyp}}(y)$, then x and y correspond to subgraphs $\Delta_{i_x}, \Delta_{i_y}$ contained in the same connected component of $(\Gamma^2)^* \setminus \mathring{B}_{\Gamma^2}(p, \nu)$. In particular, Λ' is obtained from Λ by identifying pairs of vertices, so that the natural map $\Lambda \rightarrow \Lambda'$ is π_1 -injective. Also, Λ' is a subgraph of Γ_{hyp}^2 , so that the conclusion follows from Fact 5.1. \square

We should remark that the proof of the preceding result appearing in [DS11a] uses the nontrivial fact that every closed, connected subspace of $|\Gamma|$ is arc-connected; our proof avoids using this nontrivial result and is essentially an application of the transfer principle and our construction.

Notice that the proposition implies, as in [DS11a, Corollary 18], that $\pi_1(|\Gamma|)$ is free if and only if every end has a contractible neighborhood (such ends are called *trivial*). Namely, if all ends are trivial, then $|\Gamma|$ is homotopy equivalent to a finite graph, whereas if $|\Gamma|$ contains a non-trivial end, then it contains a subgraph Γ' with exactly one non-trivial end. The compactification of Γ' is homotopy equivalent to the Hawaiian Earring and hence its fundamental group, which is contained in $\pi_1(|\Gamma|)$, is not free [CC00, Hig52].

5.2. Homology. Recall that the first homology group (with coefficients in \mathbb{Z}) $H_1(G)$ of the group G is its abelianization, i.e. the quotient of G by the subgroup $[G, G]$ generated by commutators of elements of G . The *commutator length* of $g \in G$ is defined to be infinite if $g \notin [G, G]$, while it is the minimal integer n so that g can be written as the product of n commutators if $g \in [G, G]$.

As usual, let Γ be a connected locally finite graph, and let $\theta : |\Gamma|^* \rightarrow \Gamma_{\text{hyp}}$ be the map we defined in Section 3.

Lemma 5.4. *If the loop α is null-homologous, i.e. it represents $0 \in H_1(\pi_1(|\Gamma|))$, then $\theta(\alpha)$ has finite commutator length as an element of $\pi_1(\Gamma_{\text{hyp}})$.*

Proof. Let $g \in \pi_1(|\Gamma|)$ be the element represented by α and let $f : G \rightarrow H_1(G)$ be the natural map. If $f(g) = 0$, then we can write g as the product

of, say, n commutators. As $\Theta : \pi_1(|\Gamma|) \rightarrow \pi_1(\Gamma_{\text{hyp}})$ is a group homomorphism, we have that $\Theta(g)$ can be written as a product of n commutators as well. \square

The lemma can be used to find “unexpected” non-null-homologous loops, as the commutator length of elements of a free group can be computed, see [GT79] and below. For example, it is shown in [DS10, Section 6] that, among others, the following loop α in the infinite ladder, see Figure 4 below, is not null-homologous (this loop is the same as the loop ρ as in [DS12, Figure 5.3]). The reason why this is interesting, as explained in the introduction, is that the fact that this loop is not null-homologous implies that the usual singular homology for $|\Gamma|$ is different from the topological cycle space of $|\Gamma|$.




FIGURE 4. The three parts of α

α is based at v_0 , travels through $e_1, e_2 \dots$ until it reaches the end, comes back at the top traveling in T , then loops around the first square in the clockwise direction, moves in T , loops around the second square and so on until it reaches the end. Then, it comes back at the top of the ladder staying in T . If the ladder was finite, an analogous loop would be null-homologous as, for each i , the sum of the exponents of the occurrences of e_i is 0.

Let us show that α is not null-homologous using Lemma 5.4. First, let us recall how commutators length can be computed [GT79]. Let w be a word on the letters $x_1^{\pm 1}, \dots, x_n^{\pm 1}$ representing an element in the commutator subgroup of the free group with free basis x_1, \dots, x_n . The sum of the exponents of each letter is 0, so that we can choose a *pairing* \mathcal{P} of occurrences of letters with opposite exponent. The *circle graph* associated to such a pairing is a circle C with vertices for each letter of w and appearing in the same cyclic order, together with edges $\{e_i\}$ joining paired letters.

We can associate to such a graph a matrix $M_{\mathcal{P}}$ whose ij -th entry is 1 if e_i and e_j are linked, meaning that the endpoints of e_i lie in different connected components of $C \setminus \partial e_j$, and is 0 otherwise (diagonal elements are 0). It turns out that the commutator length of w is the minimum, over all pairings \mathcal{P} , of $\text{rank}(M_{\mathcal{P}})/2$, where $M_{\mathcal{P}}$ is regarded as a matrix with coefficients in $\mathbb{Z}/2\mathbb{Z}$.

If we choose the spanning tree for Γ_{hyp} as in Figure 3, then $\theta(\alpha)$ represents the word $e_1 \dots e_\nu e_1^{-1} \dots e_\nu^{-1}$ (notice in particular that $\theta(\alpha)$ is internally null-homologous, as for each i the sum of the exponents of the occurrences of e_i is




FIGURE 5. Circle graph associated to the unique pairing for $w = x_1 x_2 x_3 x_1^{-1} x_2^{-1} x_3^{-1}$

0). It follows from the discussion above, in view of the fact that there is only one possible pairing, that the commutator length of such a word coincides with $\text{rank}(M_\nu)/2$, where M_ν is the $\nu \times \nu$ matrix whose diagonal elements are 0 and whose non-diagonal elements are 1. It is easy to show inductively that, for $n \in \mathbb{N}$, the matrix M_n defined analogously to M_ν , regarded as a matrix with coefficients in \mathbb{Z} , has determinant $(-1)^{n-1} \cdot (n-1)$ (subtracting from the first column of M_n the sum of the other columns divided by $n-2$ yields a matrix where the only non-zero element in the first column is the top entry and is $-(n-1)/(n-2)$). Consequently, if n is even, then the rank of M_n is n . If n is odd, then since M_{n-1} is a minor of M_n , we have that the rank of M_n is $n-1$. Thus, the commutator length of $e_1 \cdots e_\nu e_1^{-1} \cdots e_\nu^{-1}$ is $\text{rank}(M_\nu)/2 = \lfloor \nu/2 \rfloor > \mathbb{N}$ and, by Lemma 5.4, α is not null-homologous.

Remarks 5.5.

- (1) There are two versions of the topological cycle space, one with coefficients in $\mathbb{Z}/2\mathbb{Z}$, the other one with coefficients in \mathbb{Z} . According to [Die10a][DS10, Theorem 5], a loop α gives rise to the trivial element of the topological cycle space if and only if it traverses every edge an even number of times in the first case, or it traverses every edge the same number of times in both directions in the second case. It is readily seen that this happens if and only if $\theta(\alpha)$ is trivial in $H_1(\Gamma_{\text{hyp}}, \mathbb{Z}/2\mathbb{Z})$ in the first case, and in $H_1(\Gamma_{\text{hyp}}, \mathbb{Z})$ in the second case.
- (2) Observe that in the above argument, we did not use that Θ was injective. This fact is compatible with the standard proof that α is not nullhomologous appearing in [DS10] as the injectivity of the corresponding map $\pi_1(|\Gamma|) \rightarrow F_\infty$ was not used either.
- (3) It is interesting to note that for any loop $\alpha : [0, 1] \rightarrow |\Gamma|$, if $\Theta(\alpha)$ is internally nullhomotopic, then α is nullhomotopic, while the corresponding statement for homology is not true, as witnessed by the above discussion.
- (4) In [Die10b], a proof is given that the finite version of the loop α discussed above (called ρ there) is nullhomologous. Their proof for the finite version is more topological than our algebraic proof and they remark “But we cannot imitate this proof for ρ and our infinite ladder L , because homology classes in $H_1(|G|)$ are still finite chains: we cannot add infinitely many boundaries to subdivide ρ infinitely often.” In our opinion, this kind of quote is exactly the kind of thought process that makes nonstandard methods so powerful.

REFERENCES

- [BB11] Eli Berger and Henning Bruhn. Eulerian edge sets in locally finite graphs. *Combinatorica*, 31(1):21–38, 2011.
- [CC00] J. W. Cannon and G. R. Conner. The combinatorial structure of the Hawaiian earring group. *Topology Appl.*, 106(3):225–271, 2000.
- [CG05] Christophe Champetier and Vincent Guirardel. Limit groups as limits of free groups. *Israel J. Math.*, 146:1–75, 2005.
- [Dav77] Martin Davis. Applied nonstandard analysis. pages xii+181, 1977. Pure and Applied Mathematics.
- [Die10a] Reinhard Diestel. Graph theory. 173:xviii+437, 2010.
- [Die10b] Reinhard Diestel. Locally finite graphs with ends: A topological approach, II. Applications. *Discrete Math.*, 310(20):2750 – 2765, 2010.
- [Die11] Reinhard Diestel. Locally finite graphs with ends: a topological approach, I. Basic theory. *Discrete Math.*, 311(15):1423–1447, 2011.
- [DK04] Reinhard Diestel and Daniela Kühn. On infinite cycles. I. *Combinatorica*, 24(1):69–89, 2004.
- [DS10] Reinhard Diestel and Philipp Sprüssel. The homology of a locally finite graph with ends. *Combinatorica*, 30(6):681–714, 2010.
- [DS11a] Reinhard Diestel and Philipp Sprüssel. The fundamental group of a locally finite graph with ends. *Adv. Math.*, 226(3):2643–2675, 2011.

- [DS11b] Reinhard Diestel and Philipp Sprüssel. On the homology of locally compact spaces with ends. *Topology Appl.*, 158(13):1626–1639, 2011.
- [DS12] Reinhard Diestel and Philipp Sprüssel. Locally finite graphs with ends: a topological approach, III. Fundamental group and homology. *Discrete Math.*, 312(1):21–29, 2012.
- [Geo09a] Agelos Georgakopoulos. Infinite Hamilton cycles in squares of locally finite graphs. *Adv. Math.*, 220(3):670–705, 2009.
- [Geo09b] Agelos Georgakopoulos. Topological circles and Euler tours in locally finite graphs. *Electron. J. Combin.*, 16(1):Research Paper 40, 16, 2009.
- [Gol11] Isaac Goldbring. Ends of groups: a nonstandard perspective. *J. Log. Anal.*, 3:Paper 7, 28, 2011.
- [GT79] Richard Z. Goldstein and Edward C. Turner. Applications of topological graph theory to group theory. *Math. Z.*, 165(1):1–10, 1979.
- [Hen97] C. Ward Henson. Foundations of nonstandard analysis: a gentle introduction to nonstandard extensions. 493:1–49, 1997.
- [Hig52] Graham Higman. Unrestricted free products, and varieties of topological groups. *J. London Math. Soc.*, 27:73–81, 1952.

UNIVERSITY OF ILLINOIS AT CHICAGO, DEPARTMENT OF MATHEMATICS, STATISTICS,
AND COMPUTER SCIENCE, 851 S. MORGAN ST., CHICAGO, IL 60607-7045, USA

E-mail address: isaac@math.uic.edu

URL: <http://homepages.math.uic.edu/~isaac>

ETH ZÜRICH, DEPARTMENT OF MATHEMATICS, HG J 14.4, RÄMISTRASSE 101, 8092
ZÜRICH, SWITZERLAND

E-mail address: sisto@math.ethz.ch

URL: <http://www.math.ethz.ch/alsisto>