

CONJUGACY CLASSES OF SOLUTIONS TO EQUATIONS AND INEQUALITIES OVER HYPERBOLIC GROUPS

DANIEL GROVES AND HENRY WILTON

ABSTRACT. We study conjugacy classes of solutions to systems of equations and inequations over torsion-free hyperbolic groups, and describe an algorithm to recognize whether or not there are finitely many conjugacy classes of solutions to such a system. The class of immutable subgroups of hyperbolic groups is introduced, which is fundamental to the study of equations in this context. We apply our results to enumerate the immutable subgroups of a torsion-free hyperbolic group.

1. INTRODUCTION

This paper is concerned with decision problems relating to systems of equations and inequations over torsion-free hyperbolic groups, a subject with a long history. Makanin [17] proved that it is possible to decide whether a system of equations over a free group has a solution, and in [18] he extended this result to systems of equations and inequations. Razborov [21, 22] described an algorithm whose output gives a parametrization of the entire set of solutions to a system of equations over a free group. His description was further refined by Kharlampovich and Miasnikov [15].

Over torsion-free hyperbolic groups, Rips and Sela [24] proved that it is possible to decide if a system of equations has a solution. The corresponding result for systems of equations and inequations was proved independently by Sela [28] and Dahmani [8].

The study of equations over groups has recently received a lot of attention due to the solutions to Tarski's problem by Sela [27] and also by Kharlampovich and Miasnikov [16]. Our approach is inspired by Sela's results, and we use some of his structure theory.

Date: August 19, 2009.

The first author's work was supported in part by NSF Grant DMS-0504251. Much of the work was undertaken while the first author was a Taussky–Todd Instructor at Caltech.

Theorem A. *There is an algorithm that takes as input a finite presentation for a torsion-free hyperbolic group Γ and a finite system of equations and inequations with coefficients in Γ and determines whether or not the system has only finitely many solutions in Γ .*

If there are finitely many solutions, the algorithm provides a list of all the solutions.

In the case where Γ is a free group and the system has only equations, this theorem is immediate from Razborov's description of the set of solutions. Also, if Γ is free and there are inequations, it is not difficult to deduce Theorem A from Razborov's results. However, for general torsion-free hyperbolic Γ , more work is required.

When the system of equations and inequations has no coefficients, Γ naturally acts by conjugation on the set of solutions. Therefore, if Γ is non-elementary and there is at least one solution then there are infinitely many. Thus, in the case of coefficient-free equations and inequations, Theorem A follows immediately from the above-mentioned result of Sela and Dahmani.

It is therefore natural to consider the orbits of solutions under the action by conjugation. For example, in Sela's definition of limit groups, he takes a minimal solution in a conjugacy class in order to ensure that the limiting action on an \mathbb{R} -tree has no global fixed point. Conjugacy classes of solutions are thus of fundamental concern for Sela's approach to the Tarski problem (although this is only the very beginning of the story).

Our next result concerns conjugacy classes of solutions.

Theorem B. *There is an algorithm that takes as input a finite presentation for a torsion-free hyperbolic group Γ and a finite (coefficient-free) system of equations and inequations and determines whether or not the system has only finitely many conjugacy classes of solutions in Γ .*

If there are finitely many conjugacy classes of solutions, the algorithm provides a list containing exactly one solution in each conjugacy class.

Remark 1.1. *If Γ is free, when there is one solution (and no coefficients) there are infinitely many conjugacy classes of solutions. However, this may not be the case for arbitrary hyperbolic Γ . See Example 7.6.*

Although we believe it to be of independent interest, one of our motivations for proving Theorem B is its use in the study of Γ -limit groups (see Section 3), where Γ remains a torsion-free hyperbolic group. Unlike the case of limit groups (when Γ is free), Γ -limit groups may not be

finitely presentable. In fact, every Γ -limit group is finitely presentable if and only if Γ is coherent.

The basic building blocks of Γ -limit groups are free groups and finitely generated subgroups of Γ that admit only finitely many conjugacy classes of embeddings into Γ . We call such subgroups *immutable*. Any study of Γ -limit groups, or even the subgroups of Γ , must take into account immutable subgroups. They are one of the key technical differences between ordinary limit groups over free groups and Γ -limit groups (see Remark 7.3).

The expert reader may be interested to learn that, by the results of [28], there is an analogue over torsion-free hyperbolic groups of the analysis lattice that appears in [26]. Immutable subgroups appear at the lowest level of the analysis lattice. This analogy will be made explicit in forthcoming work [13].

We will see in Section 7 that there is a natural relationship between immutable subgroups and systems of equations and inequations with finitely many conjugacy classes of solutions (see Example 7.6). We exploit Theorem B to enumerate the immutable subgroups of Γ .

Theorem C. *Let Γ be a torsion-free hyperbolic group. The set of immutable subgroups of Γ is recursively enumerable, in the following sense: there exists a Turing machine that lists all finite subsets of Γ that generate immutable subgroups.*

This result is perhaps surprising because the finitely generated subgroups of (torsion-free) hyperbolic groups are notoriously badly behaved when it comes to decision problems. However, one of the themes of our work is that when it comes to equations and inequations, most of these difficulties can be surmounted. In [13], we will combine Theorem C with the techniques of [12] to recursively enumerate all Γ -limit groups.

Remark 1.2. *All of the algorithms we provide are uniform, in the sense that they can take as input a finite presentation for Γ as well as the equations and inequations. Thus, they work for all torsion-free hyperbolic groups in a uniform manner.*

An outline of this paper is as follows. In Section 2 we record the obvious rephrasing of solutions to equations and inequations over Γ in terms of homomorphisms to Γ . Section 3 contains some basic definitions and facts about Γ -limit groups. In Section 4 we introduce the notion of Γ -approximations, which are a key tool to circumvent the fact that Γ -limit groups may not be finitely presentable. In Section 5 we present the halves of the algorithms required for Theorems A and B

which recognize when the appropriate sets are infinite, whereas in Section 6 we provide the algorithms which recognize when they are finite. Finally, in Section 7 we discuss immutable subgroups of Γ and prove Theorem C. We end with some questions about immutable subgroups.

2. EQUATIONS AND HOMOMORPHISMS

The class of (word-)hyperbolic groups is well known—for an introduction, see for instance [5, Part III]. Throughout this paper Γ will be a fixed non-elementary torsion-free hyperbolic group. (The case in which Γ is elementary is interesting, but classical.) Fix a finite generating set $\mathcal{A} = \{a_1, \dots, a_k\}$ for Γ . Let $F(\mathcal{A})$ be the free group on \mathcal{A} , and $\theta: F(\mathcal{A}) \rightarrow \Gamma$ the canonical epimorphism. We will blur the distinction between the a_i as elements of $F(\mathcal{A})$ and as elements of Γ .

Let $X = \{x_1, \dots, x_n\}$ be a finite set of variables, and $F(X)$ the free group on X .

A *system of equations* is a collection $\Sigma \subseteq F(X) * F(\mathcal{A})$. Write $\Sigma = \{\sigma_i(\underline{x}, \underline{a})\}_{i \in I}$. A *solution* to Σ is a collection $\{\gamma_k\}_{k=1}^n \subseteq \Gamma$ so that for each $i \in I$ we have $\sigma_i(\underline{\gamma}, \underline{a}) =_{\Gamma} 1$. In other words, the result of substituting γ_j for x_j in each σ_i (and interpreting the a_k as the fixed generators for Γ) yields the trivial element of Γ .

A *system of equations and inequations* is a pair of subsets $\Sigma, \Lambda \subseteq F(X) * F(\mathcal{A})$. Suppose that $\Lambda = \{\lambda_j(\underline{x}, \underline{a})\}_{j \in J}$. A *solution* to $\Sigma = 1$, $\Lambda \neq 1$ is a subset $\underline{\gamma} \subseteq \Gamma$ so that for each $i \in I$ we have $\sigma_i(\underline{\gamma}, \underline{a}) = 1$ and for each $j \in J$ we have $\lambda_j(\underline{\gamma}, \underline{a}) \neq 1$.

The letters a_i (and a_i^{-1}) appearing in elements of Σ and Λ are called *coefficients*. We say that Σ and Λ are *missing coefficients* if they are subsets of the canonical copy of $F(X)$ in $F(X) * F(\mathcal{A})$. In case Σ and Λ are missing coefficients there are analogous subsets (still denoted Σ and Λ) of $F(X)$, which we consider to be a *coefficient-free system of equations and inequations* (also called a *system without coefficients*). Therefore, we make a distinction between systems of equations and inequations which happen not to contain any coefficients, and those which are coefficient-free. This distinction is made because in the coefficient-free setting Γ acts by conjugation on the set of solutions, whereas it does not when there are coefficients.

Following Razborov [21], it is natural to interpret solutions to systems of equations and inequations as homomorphisms to Γ . If Σ is coefficient-free, define the group H_{Σ} by the presentation

$$(1) \quad H_{\Sigma} = \langle X \mid \sigma_i(\underline{x}) = 1, i \in I \rangle.$$

If Σ has coefficients, define the group H_Σ by the presentation

$$(2) \quad \langle X, \mathcal{A} \mid \sigma_i(x, \underline{a}) = 1, i \in I \rangle.$$

In this case, the *coefficient subgroup* of H_Σ is the subgroup generated by $\mathcal{A} \subseteq H_\Sigma$.

When Σ has coefficients, the set of solutions to Σ is naturally in one-to-one correspondence with the set of homomorphisms:

$$\text{Hom}_\Gamma(H_\Sigma, \Gamma) = \{h: H_\Sigma \rightarrow \Gamma \mid h(a_i) = a_i\}.$$

If Σ is coefficient-free the set of solutions is parameterized by the full set of homomorphisms from H_Σ to Γ , denoted $\text{Hom}(H_\Sigma, \Gamma)$.

The set of solutions to $\Sigma = 1, \Lambda \neq 1$ is naturally in one-to-one correspondence with a smaller set of homomorphisms:

$$\text{IHom}_\Gamma(H_\Sigma, \Gamma; \Lambda) = \{h: H_\Sigma \rightarrow \Gamma \mid h(a_i) = a_i, h(\lambda_j) \neq 1\}.$$

(In the coefficient-free case, there is an obvious analogous definition for $\text{IHom}(H_\Sigma, \Gamma; \Lambda)$.)

Under this simple reformulation, if Σ and Λ are finite, Theorem A asserts that there is an algorithm that decides if $\text{IHom}_\Gamma(H_\Sigma, \Gamma; \Lambda)$ is finite.

If Σ and Λ are coefficient-free, we can post-compose an element of $\text{IHom}(H_\Sigma, \Gamma; \Lambda)$ by an element of $\text{Aut}(\Gamma)$, and we will get another element of $\text{IHom}(H_\Sigma, \Gamma; \Lambda)$. In particular, Γ acts on $\text{IHom}(H_\Sigma, \Gamma; \Lambda)$ by conjugation.

We denote the set $\text{IHom}(H_\Sigma, \Gamma; \Lambda)/\Gamma$ by $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$. Theorem B asserts that there is an algorithm that recognizes whether or not $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$ is finite.

3. Γ -LIMIT GROUPS

One of our key tools is the theory of Γ -*limit groups*, as introduced and studied when Γ is torsion-free hyperbolic by Sela in [28]. One of the main difficulties in studying Γ -limit groups as opposed to ordinary limit groups (when Γ is free) is that if Γ is not coherent then Γ -limit groups will not all be finitely presented.

Remark 3.1. *The Rips construction [23] can be used to construct many torsion-free hyperbolic groups which have finitely generated subgroups that are not finitely presented.*

Recall the following two definitions, due to Bestvina and Feighn in the case of free groups.

Definition 3.2. [4, Definition 1.5] *Let G be a finitely generated group, and let $\{h_n : G \rightarrow \Gamma\}$ be a sequence of homomorphisms. The stable kernel of $\{h_n\}$, denoted $\underline{Ker}(h_n)$, is the set of all elements $g \in G$ so that $g \in \ker(h_n)$ for all but finitely many n .*

The sequence $\{h_n\}$ is stable if for all $g \in G$, either $g \in \underline{Ker}(h_n)$ or, for all but finitely many n , $g \notin \ker(h_n)$.

Definition 3.3. [4, Definition 1.5] *A Γ -limit group is a group of the form $L = H_\Sigma / \underline{Ker}(h_n)$ where Σ is a coefficient-free system of equations, H_Σ is as in the previous section, and $\{h_n\} \subseteq \text{Hom}(H_\Sigma, \Gamma)$ is a stable sequence.*

A restricted Γ -limit group is a group of the form $L = H_\Sigma / \underline{Ker}(h_n)$ where Σ is a system of equations with coefficients, and $\{h_n\} \subseteq \text{Hom}_\Gamma(H_\Sigma, \Gamma)$ is a stable sequence. The image of the coefficient subgroup of H_Σ in L is called the coefficient subgroup of L .

It should be noted that restricted Γ -limit groups are, in particular, Γ -limit groups. Almost all of the theory of restricted Γ -limit groups is identical to that of Γ -limit groups. Often in this paper we will omit the adjective ‘restricted’; this should not cause the reader any confusion.

Remark 3.4. *By taking a constant sequence of homomorphisms, we see that all finitely generated subgroups of Γ are Γ -limit groups.*

Definition 3.5. *A Γ -limit group is called strict if it can be represented as $H_\Sigma / \underline{Ker}(h_n)$ where $\{h_n : H_\Sigma \rightarrow \Gamma\}$ are pairwise non-conjugate.*

A restricted Γ -limit group is strict if it can be represented as $H_\Sigma / \underline{Ker}(h_n)$ where $\{h_n : H_\Sigma \rightarrow \Gamma\}$ are pairwise distinct.

We refer the reader to [28] for the basic theory of Γ -limit groups where Γ is a torsion-free hyperbolic group. However, we do emphasize two features. Recall that if H and G are groups then we say that G is *fully residually H* if for every finite subset $A \subseteq G$ there is a homomorphism $h_A : G \rightarrow H$ which is injective on A .

Proposition 3.6. [28, Proposition 1.18] *Let Γ be a non-elementary torsion-free hyperbolic group. A finitely generated group G is fully residually Γ if and only if it is a Γ -limit group.*

The above property will be used several times in this paper without comment.

A strict Γ -limit group L is naturally equipped with a faithful action on an \mathbb{R} -tree. If L is a strict restricted Γ -limit group, the coefficient subgroup of L acts elliptically on this tree. In order to apply the

Rips Machine [25, 14] and deduce that L splits (relative to the coefficient subgroup), we need to know that the action has no global fixed point. This is the content of the following standard proposition. In the coefficient-free case, this is proved in [20], [2] or [6], for example. The case with coefficients is well-known to the experts, but as far as we are aware there is no proof in the literature. We take this opportunity to record a proof, which was communicated to the first author by Sela.

Proposition 3.7. *Any strict Γ -limit group L acts without a global fixed point on an \mathbb{R} -tree. If L is restricted the action may be chosen so that the coefficient subgroup fixes a point.*

Proof (Case with coefficients). We start by outlining the standard construction of the \mathbb{R} -tree as a pointed Gromov-Hausdorff limit; see [6] or [3], for instance, for details. Fix a word metric d on Γ . Let $\{h_n\} \subseteq \text{Hom}_\Gamma(H_\Sigma, \Gamma)$ be a stable sequence of pairwise distinct homomorphisms. For each n define the scaling factors δ_n and the basepoint $*_n = 1 \in \Gamma$. For each n we have an action of H_Σ on Γ , equipped with the rescaled word metric d/δ_n . A subsequence of this sequence of actions converges (in the pointed Gromov-Hausdorff topology) to an action of H_Σ on a Gromov-hyperbolic space T equipped with a basepoint $*$. Because the sequence $\{h_n\}$ is pairwise distinct, the scaling factors δ_n tend to infinity, and it follows that T is an \mathbb{R} -tree. The kernel of the action of H_Σ on T clearly contains $\underline{\text{Ker}}(h_n)$, so the action of H_Σ descends to an action of L on T .

For each element g of the coefficient subgroup of H_Σ , $h_n(g) = g$ for each n and so $d(1, h_n(g))/\delta_n$ tends to 0. It follows that $g* = *$, so the coefficient subgroups of H_Σ and hence L act elliptically, as required.

It remains to prove that the action of L on T has no global fixed point. This is the part of the proof that differs significantly from the coefficient-free case. The problem is that, because of the coefficient subgroup, we could not conjugate the homomorphisms h_n to make $*_n$ centrally located. Instead, we argue as follows.

The proof of [28, Lemma 1.3(vi)] applies to show that segment stabilizers of the action of L on T are abelian. The coefficient subgroup of L is clearly isomorphic to Γ so is non-abelian. Therefore, the basepoint $*$ is the unique point fixed by the coefficient subgroup. On the other hand, by construction some generator x_i moves $*$ distance 1, so there is no global fixed point. \square

We are particularly interested in cyclic splittings that give rise to infinitely many outer automorphisms.

Definition 3.8. *A cyclic splitting is essential if the edge group is of infinite index in any vertex group.*

The point of Theorem 3.9 is that one can choose the splitting assured by the Rips Machine to be essential.

Theorem 3.9 (Sela, [28]). *Let Γ be a non-elementary torsion-free hyperbolic group, and let L be a non-abelian, freely indecomposable, strict Γ -limit group. Then L admits an essential cyclic splitting.*

If, furthermore, L is a restricted Γ -limit group then the splitting may be chosen so that the coefficient subgroup is elliptic.

Sela needs Theorem 3.9 in [28, Section 2] (and claims that it is true on page 8). However, he does not prove it in [28]. In [26, Theorem 3.2] it is proved when Γ is free. There are only a couple of extra subtleties when Γ is not free, and essentially the same proof works (in particular, we know that Sela has a proof of Theorem 3.9). However, we feel uncomfortable simply asserting that this strategy works so we include our own proof here. The proof we present is inspired by a remark towards the beginning of the proof of [28, Theorem 3.2], that Theorem 3.9 is straightforward once one knows that all abelian subgroups of L are finitely generated.

Proof of Theorem 3.9. Note that Γ is a toral relatively hyperbolic group. Therefore we may apply the results of [11] to L . The advantage of this approach is that Theorem 3.9 is false in this context (see [11, Remark 5.1]). Therefore, the proof of [11, Corollary 5.12] does not rely on the result we are currently trying to prove, so there is no risk of circularity. The content of [11, Corollary 5.12] is that abelian subgroups of L are finitely generated.

Consider the abelian JSJ decomposition Δ of L and collapse all edges that give rise to non-essential cyclic splittings of L , yielding a graph of groups Δ' for L . As pointed out in the proof of [26, Theorem 3.2], the action on an \mathbb{R} -tree that arises from Proposition 3.7 leads (via [28, Theorem 1.5]) to either an essential cyclic splitting or a non-trivial abelian splitting of L . Therefore, we may assume that Δ' is non-trivial.

Suppose that Δ' contains only non-cyclic edge groups, and let V be a non-abelian vertex group of this decomposition. Then V contains a non-cyclic abelian subgroup, so cannot be a subgroup of Γ . Hence V is a strict Γ -limit group, and so admits a non-trivial abelian splitting Θ in which all non-cyclic abelian subgroups are elliptic, by [28, Lemma 1.7]. As noted above, Θ cannot be an inessential cyclic splitting. The decomposition Θ of V can be used to refine Δ' . If Θ has a trivial edge group, then we obtain a contradiction to the fact that L is freely

indecomposable. Otherwise, V is obtained from a rigid vertex group W of Δ by adjoining finitely many roots and each such root is elliptic in every abelian splitting of L . Therefore, Θ can be used to refine Δ , and we obtain an abelian splitting of L in which W is not elliptic, contradicting the canonical properties of Δ .

In the case with coefficients, the proof is almost identical. The only modification is that one works with the *restricted* abelian JSJ decomposition, rather than the full abelian JSJ decomposition. The restricted decomposition is obtained by only considering those splittings in which the coefficient subgroup is elliptic. \square

4. Γ -APPROXIMATIONS AND SPLITTINGS

We have already noted that Γ -limit groups need not be finitely presented. Implementing algorithms involving groups that are not finitely presented is obviously more difficult than when the groups are finitely presented. We will circumvent this difficulty by working with finitely presented approximations to our Γ -limit groups.

Remark 4.1. *Any homomorphism $\phi: G \rightarrow H$ induces a map*

$$\phi^*: \text{Hom}(H, \Gamma) \rightarrow \text{Hom}(G, \Gamma)$$

via $\phi^(h) = h \circ \phi$. If ϕ is an epimorphism then ϕ^* is an injection.*

Definition 4.2. *Suppose that G is a finitely generated group. A Γ -approximation to G is a finitely presented group \hat{G} together with an epimorphism $\eta: \hat{G} \rightarrow G$ so that η^* is a bijection.*

Often the map η will be implicit, and we will simply refer to \hat{G} as a Γ -approximation to G .

We also need a version of Γ -approximations for equations with coefficients.

Definition 4.3. *Let Σ, Σ_0 be systems of equations with coefficients, and H_Σ, H_{Σ_0} the associated groups. If $\Sigma_0 \subseteq \Sigma$ then there is a natural epimorphism $\eta: H_{\Sigma_0} \rightarrow H_\Sigma$ which in turn induces an injection*

$$\eta^*: \text{Hom}_\Gamma(H_\Sigma, \Gamma) \hookrightarrow \text{Hom}_\Gamma(H_{\Sigma_0}, \Gamma).$$

A Γ -approximation to H_Σ is a pair (H_{Σ_0}, η) where Σ_0 is a finite system of equations and η^ is a bijection.*

Saying that η^* is a bijection is equivalent to saying that every homomorphism from \hat{G} to Γ factors through η . To prove the existence of approximations, we appeal to the following result from [28], which states that torsion-free hyperbolic groups are *equationally Noetherian*.

Theorem 4.4. [28, Theorem 1.22] *Let Γ be a torsion-free hyperbolic group. Then every infinite system of equations in finitely many variables over Γ is equivalent to a finite subsystem.*

Note that the above result is only proved in [28] for equations with no coefficients. However, the more general result may be proved with exactly the same proof using the restricted Makanin–Razborov diagram rather than the (ordinary) Makanin–Razborov diagram.

Translating Theorem 4.4 into the language of homomorphisms ensures that approximations will always exist.

Theorem 4.5. *Suppose that Γ is a torsion-free hyperbolic group. For any sequence of epimorphisms of groups*

$$G_1 \rightarrow G_2 \rightarrow G_3 \rightarrow \dots$$

the corresponding sequence of homomorphism varieties

$$\mathrm{Hom}(G_1, \Gamma) \supset \mathrm{Hom}(G_2, \Gamma) \supset \mathrm{Hom}(G_3, \Gamma) \dots$$

eventually stabilizes.

In particular, for any finitely generated group G there is a finitely presented group \hat{G} along with an epimorphism $\eta: \hat{G} \rightarrow G$ so that the injection $\eta^: \mathrm{Hom}(G, \Gamma) \rightarrow \mathrm{Hom}(\hat{G}, \Gamma)$ is bijective.*

If G has a cyclic splitting, we can choose a Γ -approximation \hat{G} that admits a splitting of the same form. To do this, we give a recipe for building a Γ -approximation \hat{G} for a group G , given a graph-of-groups decomposition Λ of G with cyclic edge groups, and Γ -approximations for the vertex groups of Λ .

Lemma 4.6. *Suppose $G = A *_C B$ is a splitting of G over a (possibly trivial) cyclic subgroup C . Let $\eta_A: \hat{A} \rightarrow A$ and $\eta_B: \hat{B} \rightarrow B$ be Γ -approximations and let $\hat{C}_A \subseteq \hat{A}$ and $\hat{C}_B \subseteq \hat{B}$ be (possibly trivial) cyclic subgroups with the property that $\eta_A(\hat{C}_A) = \eta_B(\hat{C}_B) = C$. Then*

$$\hat{G} = \hat{A} *_{\hat{C}_A = \hat{C}_B} \hat{B}$$

is a Γ -approximation to G .

Proof. We can characterize $\mathrm{Hom}(G, \Gamma)$ as

$$\mathrm{Hom}(G, \Gamma) = \{(\alpha, \beta) \in \mathrm{Hom}(A, \Gamma) \times \mathrm{Hom}(B, \Gamma) \mid \alpha|_C = \beta|_C\}.$$

We can characterize $\mathrm{Hom}(\hat{G}, \Gamma)$ similarly, and the result follows immediately. \square

Likewise, we have a similar result for HNN-extensions.

Lemma 4.7. *Suppose $G = A *_C \sim_{C'}$ is a splitting of G , where the stable letter conjugates the infinite cyclic subgroups C and C' . Let $\eta_A : \hat{A} \rightarrow A$ be a Γ -approximation and let $\hat{C} \subseteq \hat{A}$ and $\hat{C}' \subseteq \hat{A}$ be cyclic subgroups with the property that $\eta_A(\hat{C}) = C$ and $\eta_A(\hat{C}') = C'$. Then*

$$\hat{G} = \hat{A} *_C \sim_{\hat{C}'}$$

is a Γ -approximation to G .

Proof. We can characterize $\text{Hom}(G, \Gamma)$ as

$$\text{Hom}(G, \Gamma) = \{(\alpha, \gamma) \in \text{Hom}(A, \Gamma) \times \Gamma \mid \alpha|_C = \iota_\gamma \circ \alpha|_{C'}\}.$$

where $\iota_\gamma \in \text{Aut}(\Gamma)$ is conjugation by γ . Just as before, $\text{Hom}(\hat{G}, \Gamma)$ admits a similar characterization and the result follows. \square

5. IF THERE ARE INFINITELY MANY

5.1. A characterization in terms of splittings of quotients. In this section we discuss the cases when $\text{IHom}_\Gamma(H_\Sigma, \Gamma; \Lambda)$ and $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$ are infinite. The idea is to characterize the existence of infinitely many homomorphisms in terms of splittings of quotients of H_Σ .

Proposition 5.1. *If $\text{IHom}_\Gamma(H_\Sigma, \Gamma; \Lambda)$ is infinite then there is a strict restricted Γ -limit quotient $\eta : H_\Sigma \rightarrow L$ so that infinitely many elements of $\text{IHom}_\Gamma(H_\Sigma, \Gamma; \Lambda)$ factor through η . The quotient L admits a free or essential cyclic splitting, and the coefficient subgroup is elliptic in this splitting.*

If $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$ is infinite then there is a strict Γ -limit quotient $\eta : H_\Sigma \rightarrow L$ so that for infinitely many $o \in \mathcal{O}(H_\Sigma, \Gamma; \Lambda)$, every representative of o factors through η . The quotient L is either abelian or admits a free or essential cyclic splitting.

Proof. Consider the case with coefficients. Assuming $\text{IHom}_\Gamma(H_\Sigma, \Gamma; \Lambda)$ is infinite, let $h_n \in \text{IHom}_\Gamma(H_\Sigma, \Gamma; \Lambda)$ be a sequence of pairwise distinct homomorphisms. Passing to a subsequence, we may assume that (h_n) is a stable sequence. Let $L = H_\Sigma / \underline{\text{Ker}}(h_n)$. By [28, Theorem 1.17], we may pass to a further subsequence and assume that every h_n factors through the quotient map $\eta : H_\Sigma \rightarrow L$. Since the h_n are distinct, L is a strict restricted Γ -limit group. The quotient L is non-abelian and so by Theorem 3.9 admits a free or essential cyclic splitting in which the coefficient subgroup is elliptic.

In the coefficient-free case, the fact that $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$ is infinite implies that there exists a sequence of pairwise non-conjugate homomorphisms $h_n : H_\Sigma \rightarrow \Gamma$. As in the previous paragraph, we can pass to a stable subsequence such that every h_n factors through $\eta : H_\Sigma \rightarrow L =$

$H_\Sigma/\overrightarrow{\text{Ker}}(h_m)$. Once again, Theorem 3.9 ensures that if L is non-abelian then L splits as required. \square

In fact, Proposition 5.1 has a strong converse. A simple version of this converse is encapsulated in the following two results.

Proposition 5.2. *Suppose that Σ and Λ have coefficients. Suppose H_Σ admits a strict Γ -limit quotient $\eta : H_\Sigma \rightarrow L$ so that*

- (1) *some element of $\text{IHom}_\Gamma(H_\Sigma, \Gamma; \Lambda)$ factors through η ; and*
- (2) *L admits a free or essential cyclic splitting in which the coefficient subgroup is elliptic.*

Then $\text{IHom}_\Gamma(H_\Sigma, \Gamma; \Lambda)$ is infinite.

Proposition 5.3. *Suppose that Σ and Λ are coefficient-free. Suppose H_Σ admits a strict Γ -limit quotient $\eta : H_\Sigma \rightarrow L$ for which*

- (1) *a representative of some element of $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$ factors through η ; and*
- (2) *L is abelian or admits a free or essential cyclic splitting.*

Then $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$ is infinite.

We will prove stronger results than these, which are more easily applied in the algorithmic context (Propositions 5.12 and 5.13). However, for the moment it is useful to think about how one might go about proving these results. The strategy is to pre-compose by Dehn twists, generalized Dehn twists or partial conjugations (see Definition 5.11) to explicitly exhibit infinitely many elements of $\text{IHom}_\Gamma(H_\Sigma, \Gamma; \Lambda)$ or $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$. If we are considering only equations, this is straightforward. In general, however, we need to ensure that the inequalities are preserved when twisting by non-inner automorphisms of L .

In order to do this, we use the following result. It is well known in the case where Γ is a free group, and is often called Baumslag's Lemma. (Note that our proof can be substantially simplified in the free case.)

Lemma 5.4. *Suppose $z, a_1, \dots, a_n \in \Gamma$ and $[a_i, z] \neq 1$ for all i . Then whenever $|m_i|$ are sufficiently large for $i < n$, the word*

$$w = a_1 z^{m_1} \dots a_n z^{m_n}$$

does not commute with z .

The proof uses the δ -hyperbolic geometry of the Cayley graph of Γ . Let $d_{\mathcal{A}}$ be the word metric on the Cayley graph of Γ with respect to \mathcal{A} and suppose this metric is δ -hyperbolic. To prove Lemma 5.4 we need a version of the Margulis Lemma for hyperbolic groups. The following is essentially contained in [20].

Lemma 5.5. *Let $a, b, c, d \in \Gamma$ and suppose that c^a does not commute with d^b . If m and n are integers, write $I_m = a[1, c^m]$ and $J_n = b[1, d^n]$ for choices of geodesic arcs in the Cayley graph of Γ . Fix $\epsilon \geq 0$. Then the diameter of*

$$I_m \cap N_\epsilon(J_n)$$

(where $N_\epsilon(J_n)$ denotes the ϵ -neighbourhood of J_n) is uniformly bounded, independently of m and n .

Proof. Let C_0 be the number of elements in the ball of radius $8\delta + 4\epsilon$. Let $j \leq C_0$.

Let $I' \subseteq I_m$ be the arc of points at distance at least $|c| + 2\delta$ from the endpoints of I_m . Then for $x \in I'$, $d_{\mathcal{A}}((c^a)^i x, I_m) \leq 2\delta$. Let π be (a choice of) closest-point projection to I_m . Then $\pi \circ (c^a)^i$ is at distance at most 2δ from an orientation-preserving isometry $c' : I' \rightarrow I_m$.

Similarly, let $J' \subseteq I_m$ be an arc of points at distance at least $|d^j| + 2\delta + 2\epsilon$ from the endpoints of I_m , and also within ϵ of J_n . Then $\pi \circ (d^b)^j$ is at distance at most $2\delta + 2\epsilon$ from an isometry $d'_j : J' \rightarrow I_m$. Note that the translation length of d'_j is bounded above by a linear function of j .

For a contradiction, suppose that the diameter of $I_m \cap N_\epsilon(J_n)$ is large enough that there exists $x \in I' \cap J'$ such that $c'(x) \in J'$ and $d'_j(x) \in I'$ for all $j \leq C_0$. In particular, $d'_j c'(x) = c' d'_j(x) \in I_m$; we have that $d_{\mathcal{A}}((d^b)^j c^a(x), d'_j c'(x)) \leq 4\delta + 2\epsilon$ and, likewise, $d_{\mathcal{A}}(c^a (d^b)^j(x), c' d'_j(x)) \leq 4\delta + 2\epsilon$. It follows that $d_{\mathcal{A}}((d^b)^{-j} (c^a)^{-1} (d^b)^j c^a(x), x) \leq 8\delta + 4\epsilon$ and so, by the freeness of the action on the Cayley graph

$$(d^b)^{-j} (c^a)^{-1} (d^b)^j c^a = (d^b)^{-j'} (c^a)^{-1} (d^b)^{j'} c^a$$

for some distinct $j, j' \leq C_0$. Using the fact that torsion-free hyperbolic groups are commutative transitive, it follows that c^a commutes with d^b , as required. \square

Proof of Lemma 5.4. We use the word-hyperbolic geometry of Γ . Suppose the word metric $d_{\mathcal{A}}$ on Γ corresponding to the generating set \mathcal{A} is δ -hyperbolic and denote the word length of an element γ by $|\gamma|$.

The proof is by induction on n —note that the result is immediate for $n = 1$. If w and z commute then, by modifying m_n , we can assume that $w = 1$. Fix geodesic arcs $[1, z]$ and $[1, a_i]$ for each i . Write $w_i = a_1 z^{m_1} \dots z^{m_{i-1}} a_i$ so, by induction, we may assume that w_i and z do not commute for all i . Let

$$\alpha_i = w_{i-1} z^{m_{i-1}} [1, a_i]$$

(where w_0 is taken to be 1) and

$$\zeta_i = w_i [1, z^{m_i}]$$

so, if $w = 1$ the concatenation $\alpha_1 \cdot \zeta_1 \cdot \alpha_2 \cdots \alpha_n \cdot \zeta_n$ defines a $2n$ -sided polygon in Γ .

Any such polygon is uniformly $2(n-1)\delta$ -slim, so in particular ζ_1 lies within the $2(n-1)\delta$ -neighbourhood of the union of the α_i (over all i) and ζ_i (for $i > 1$). A segment of length at most $|a_i| + 4(n-1)\delta$ can lie within the neighbourhood of α_i , so if $A = \sum_i |a_i|$ then the lengths of segments of ζ_1 covered by neighbourhoods of the other ζ_i sum to at least $|z^{m_1}| - A - 4n(n-1)\delta$. It follows easily from the fact that rectangles are 4δ -slim that, if two geodesic segments have endpoints that lie within $2(n-1)\delta$ of each other, then the entire segments lie within $2n\delta$ of each other. Hence there exists a segment of ζ_1 of length at least

$$\frac{1}{n} (|z^{m_1}| - A - 4n(n-1)\delta)$$

that lies within a $2n\delta$ -neighbourhood of some other ζ_i . Therefore by Lemma 5.5, if $|m_1|$ is sufficiently large $a_1^{-1}w_i$ commutes with z , contradicting the inductive hypothesis. \square

The following propositions are easy consequences.

Proposition 5.6. *Suppose that $a_1, \dots, a_k, b_1, \dots, b_k, c$ are elements of Γ so that $a_1 b_1 \cdots a_k b_k \neq 1$. Then for all but finitely many j , the element*

$$(c^{-j} a_1 c^j) b_1 \cdots (c^{-j} a_k c^j) b_k$$

is non-trivial in Γ .

Proposition 5.7. *Suppose that a_1, \dots, a_k, t, c are elements of Γ so that $a_1 t^{n_1} \cdots a_k t^{n_k} \neq 1$. Then for all but finitely many j , the element*

$$a_1 (tc^j)^{n_1} \cdots a_k (tc^j)^{n_k}$$

is non-trivial in Γ .

5.2. Approximate quotients. In the algorithmic setting, the problem with applying Propositions 5.2 and 5.3 is that in general the limit quotient L may not be finitely presented. We deal with this by using Γ -approximations.

Remark 5.8. *Let $\pi : H_\Sigma \rightarrow L$ be any quotient of H_Σ and let $\eta : \hat{L} \rightarrow L$ be a Γ -approximation to L . By definition $\eta^* : \text{Hom}(L, \Gamma) \rightarrow \text{Hom}(\hat{L}, \Gamma)$ is a bijection, and $\pi^* : \text{Hom}(L, \Gamma) \rightarrow \text{Hom}(H_\Sigma, \Gamma)$ is an injection. So the composition*

$$\pi^* \circ (\eta^*)^{-1} : \text{Hom}(\hat{L}, \Gamma) \rightarrow \text{Hom}(H_\Sigma, \Gamma)$$

is an injection.

In the coefficient-free case, if $\hat{L} \cong \langle X \mid \mathcal{R} \rangle$ then the sentence

$$\forall \underline{x} \ (\mathcal{R}(\underline{x}) \implies \Sigma(\underline{x}))$$

holds over Γ . Likewise, in the case with coefficients, if $\hat{L} \cong \langle X, \mathcal{A} \mid \mathcal{R} \rangle$ then the sentence

$$\forall \underline{x} \ (\mathcal{R}(\underline{x}, \underline{a}) \implies \Sigma(\underline{x}, \underline{a}))$$

holds in Γ .

Therefore, in order to show that $\text{IHom}_\Gamma(H_\Sigma, \Gamma; \Lambda)$ and $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$ are infinite, we will show that there are infinitely many elements of $\text{IHom}_\Gamma(\hat{L}, \Gamma; \Lambda)$ and $\mathcal{O}(\hat{L}, \Lambda, \Gamma)$.

5.3. Dehn twists, generalized Dehn twists and partial conjugations. In this subsection we explain why the fact that the group \hat{L} has a splitting is enough to deduce that $\text{IHom}_\Gamma(\hat{L}, \Gamma; \Lambda)$ and $\mathcal{O}(\hat{L}, \Gamma; \Lambda)$ are infinite.

Definition 5.9. *Suppose that G is a finitely generated non-abelian group with a free or essential cyclic splitting Θ . A homomorphism $h : G \rightarrow \Gamma$ is called a witness with respect to Θ if*

- (1) h is injective on edge groups C ;
- (2) $h(G)$ is non-abelian;
- (3) if A is a non-abelian vertex group and the edge group C is non-trivial then $h(C)$ is non-central in $h(A)$; and
- (4) if Θ is a free splitting then the image of each vertex group and each stable letter under h is non-trivial. Furthermore, if A is a non-abelian vertex group, $h(A)$ is non-abelian.

Remark 5.10. *If G is a non-abelian strict Γ -limit group then a witness always exists.*

A witness reflects the structure of the splitting Θ well enough that Dehn twists, generalized Dehn twists and partial conjugation give new (conjugacy classes of) homomorphisms to Γ .

Definition 5.11. *There are sets of groups automorphisms associated to one-edge free or essential cyclic splittings.*

- (1) If $G = A * B$ and $g \in A$ we can define the partial conjugation $\alpha_g \in \text{Aut}(H)$ such that $\alpha_g(a) = a$ for all $a \in A$ and $\alpha_g(b) = gb g^{-1}$ for all $b \in B$.
- (2) Suppose $G = A *_C B$ and C is not central in B . If $C = \langle c \rangle$ then we can define the Dehn twist $\delta_c \in \text{Aut}(G)$ by $\delta_c(a) = a$ for $a \in A$ and $\delta_c(b) = bcb^{-1}$ for each $b \in B$.

- (3) Suppose $G = A *_C B$ and B is abelian and torsion-free. A generalized Dehn twist is an automorphism $\delta' \in \text{Aut}(G)$ that acts trivially on A and that acts via a unimodular linear automorphism of B , fixing C .
- (4) Suppose $G = A *_C$ and let t be the stable letter. If $C = \langle c \rangle$ then we can define the Dehn twist $\delta_c \in \text{Aut}(G)$ by $\delta_c(a) = a$ for $a \in A$ and $\delta_c(t) = tc$.

Proposition 5.12. *Suppose that G is a group and that Θ is a non-trivial (one-edge) free or essential cyclic splitting of G . Let $X \cup \mathcal{A}$ be a finite generating set for G , and assume that the coefficient subgroup $\langle \mathcal{A} \rangle$ is elliptic in Θ . Suppose furthermore that $\Lambda(\underline{x}, \underline{a}) \neq 1$ is a finite system of inequations, and that there is a homomorphism $h : G \rightarrow \Gamma$ so that*

- (i) h is a witness for Θ ;
- (ii) the restriction of h to the coefficient subgroup $\langle \mathcal{A} \rangle$ surjects onto Γ ; and
- (iii) the inequations $\Lambda(h(\underline{x}), h(\underline{a})) \neq 1$ hold in Γ .

Then $\text{IHom}_\Gamma(G, \Gamma; \Lambda)$ is infinite.

Proof. The idea is to twist h by partial conjugations, Dehn twists or generalized Dehn twists to construct infinitely many elements of $\text{Hom}(G, \Gamma)$ and then to use Propositions 5.6 and 5.7 to conclude that all but finitely many of these elements satisfy the inequalities $\Lambda \neq 1$. There are four cases to consider.

If Θ is a free splitting $G = A * B$ then we can assume that the coefficient subgroup is contained in A . Note that there is an element $g \in A$ such that $h(g)$ does not centralize $h(B)$. The sequence of homomorphisms $(h \circ \alpha_g^n)_{n \in \mathbb{N}}$ exhibits infinitely many elements of $\text{Hom}_\Gamma(G, \Gamma)$ and, by Proposition 5.6, the inequalities $\Lambda(h \circ \alpha_g^n(\underline{x}), h(\underline{a})) \neq 1$ do not hold for at most finitely many n . Therefore $\text{IHom}_\Gamma(G, \Gamma; \Lambda)$ is infinite.

Suppose that Θ is an essential cyclic splitting of the form $G = A *_C B$ with A and B non-abelian and the coefficient subgroup is contained in A . If $C = \langle c \rangle$ then $h(c)$ does not lie in the centre of $h(B)$ so, exactly as in the case of the free product, infinitely many of the homomorphisms $(h \circ \delta_c^n)_{n \in \mathbb{N}}$ are in $\text{IHom}_\Gamma(G, \Gamma; \Lambda)$.

Suppose now that Θ is an essential cyclic splitting of the form $G = A *_C B$ with A non-abelian and B abelian, and the coefficient group is contained in A . As the splitting is essential, B is of rank at least 2. Write $B = B_0 \oplus \langle t \rangle$ where $h(t) \neq 1$ and $C \subseteq B_0$. Define the generalized Dehn twist $\delta' \in \text{Aut}(G)$ by $\delta'(a) = a$ for $a \in A$, $\delta'(b) = b$ for $b \in B_0$ and $\delta'(t) = tc$, where $c \in C \setminus \{1\}$. Note that since h is a witness,

$h \neq h \circ \delta'$. The sequence of homomorphisms $(h \circ (\delta')^n)_{n \in \mathbb{N}}$ exhibits infinitely many elements of $\text{Hom}_\Gamma(G, \Gamma)$ and, by Proposition 5.7, the inequalities $\Lambda(h \circ (\delta')^n(\underline{x}), h(\underline{a})) \neq 1$ do not hold for at most finitely many n . Therefore $\text{IHom}_\Gamma(G, \Gamma; \Lambda)$ is infinite.

Suppose finally that Θ is a cyclic splitting of the form $G = A *_C$. If $C = \langle c \rangle$ then the sequence of homomorphisms $(h \circ \delta_c^n)_{n \in \mathbb{N}}$ exhibits infinitely many elements of $\text{Hom}_\Gamma(G, \Gamma)$. As before, it follows from Proposition 5.7 that $\text{IHom}_\Gamma(G, \Gamma; \Lambda)$ is infinite. \square

We have a similar result for $\mathcal{O}(G, \Gamma; \Lambda)$.

Proposition 5.13. *Suppose that G is a non-abelian group and that Θ is a non-trivial (one-edge) free or essential cyclic splitting of G , and Θ has at least one non-abelian vertex group. Let X be a finite generating set for G . Suppose furthermore that $\Lambda(\underline{x}) \neq 1$ is a finite system of inequations, and that there is a homomorphism $h : G \rightarrow \Gamma$ so that*

- (i) h is a witness for Θ ; and
- (ii) the inequations $\Lambda(h(\underline{x})) \neq 1$ hold in Γ .

Then $\mathcal{O}(G, \Lambda, \Gamma)$ is infinite.

The proof is identical to the proof of Proposition 5.12. After noting that the centralizer of any non-abelian subgroup of Γ is trivial, it is straightforward to verify that the homomorphisms constructed are pairwise non-conjugate.

Remark 5.14. *If G is abelian and $\mathcal{O}(G, \Gamma; \Lambda)$ contains a non-trivial conjugacy class of homomorphisms then $\mathcal{O}(G, \Gamma; \Lambda)$ is infinite. Similarly, if G is a free product of abelian groups and $\mathcal{O}(G, \Gamma; \Lambda)$ is non-empty then $\mathcal{O}(G, \Gamma; \Lambda)$ is infinite.*

5.4. Checking for infinitely many solutions. We have assembled all the tools we need to check that $\text{IHom}_\Gamma(H_\Sigma, \Gamma; \Lambda)$ and $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$ are infinite.

Definition 5.15. *A presentation for a group G is said to exhibit a free splitting $A * B$ if it is of the form*

$$\langle \mathcal{A}, \mathcal{B} \mid \mathcal{R}_\mathcal{A}, \mathcal{R}_\mathcal{B} \rangle$$

where $\mathcal{R}_\mathcal{A}$ involves only generators in \mathcal{A} and $\mathcal{R}_\mathcal{B}$ involves only generators in \mathcal{B} (and $A \cong \langle \mathcal{A} \mid \mathcal{R}_\mathcal{A} \rangle$ and $B \cong \langle \mathcal{B} \mid \mathcal{R}_\mathcal{B} \rangle$).

*Likewise, a presentation for a group G is said to exhibit a cyclic splitting $A *_C B$ if it is of the form*

$$\langle \mathcal{A}, \mathcal{B}, c \mid \mathcal{R}_\mathcal{A}, \mathcal{R}_\mathcal{B}, c = w_\mathcal{A}, c = w_\mathcal{B} \rangle$$

where \mathcal{R}_A involves only generators in \mathcal{A} , \mathcal{R}_B involves only generators in \mathcal{B} , w_A is a word in the generators \mathcal{A} and w_B is a word in the generators \mathcal{B} . (Here $A \cong \langle \mathcal{A} \mid \mathcal{R}_A \rangle$, $B \cong \langle \mathcal{B} \mid \mathcal{R}_B \rangle$ and $C = \langle c \rangle$.)

There is an obvious analogous definition of what it means for a presentation to exhibit an HNN-extension over a cyclic subgroup.

In what follows, any vertex group is treated as non-abelian unless its relations include commutators of each pair of generators.

Definition 5.16. *A presentation that exhibits a cyclic splitting is said to exhibit an essential cyclic splitting if whenever a vertex group is visibly abelian it is not virtually cyclic.*

Note that it is easy to check algorithmically if a presentation exhibits a free or essential cyclic splitting, and to check whether a vertex group is visibly abelian. Of course, one cannot in general decide if the edge group of a cyclic splitting is infinite, or indeed if a splitting is non-trivial. In our algorithms we will find witnesses that prove that our splittings are indeed as they appear.

Theorem 5.17. *Let Γ be a torsion-free hyperbolic group and $\Sigma = 1, \Lambda \neq 1$ a finite system of equations and inequations. There exists a Turing machine that takes as input Σ and Λ and terminates if and only if the set of solutions $\text{IHom}_\Gamma(H_\Sigma, \Gamma; \Lambda)$ is infinite.*

Proof. The Turing machine systematically enumerates group presentations of the form

$$\hat{L} \cong \langle X, \mathcal{A} \mid \mathcal{R} \rangle$$

and checks to see whether \hat{L} satisfies the following conditions:

- (1) the presentation exhibits a free or essential cyclic splitting Θ of \hat{L} in which the coefficient subgroup is elliptic;
- (2) for each vertex group V of Θ which is not visibly abelian there is a homomorphism $f: \hat{L} \rightarrow \Gamma$ so that $f(V)$ is non-abelian;
- (3) there exists a witness h for Θ that satisfies the inequations $\Lambda(h(\underline{x}), h(\underline{a})) \neq 1$ —note that the existence of a witness can be checked by solving a finite system of equations and inequations over Γ ; and
- (4) the sentence

$$\forall \underline{x} (\mathcal{R}(\underline{x}, \underline{a}) \implies \Sigma(\underline{x}, \underline{a}))$$

holds in Γ .

If such a group \hat{L} is found then $\text{IHom}_\Gamma(\hat{L}, \Gamma; \Lambda)$ is infinite, by Proposition 5.12. But, since

$$\mathcal{R}(\underline{x}, \underline{a}) \implies \Sigma(\underline{x}, \underline{a})$$

it follows that $\text{IHom}_\Gamma(H_\Sigma, \Gamma; \Lambda)$ is infinite by Remark 5.8.

It remains to show that such an \hat{L} always exists if $\text{IHom}_\Gamma(H_\Sigma, \Gamma; \Lambda)$ is infinite. By Proposition 5.1, H_Σ has a limit group quotient L with a free or essential cyclic splitting in which the coefficient group is elliptic. By Lemmas 4.6 and 4.7, there exists a Γ -approximation

$$\hat{\eta} : \hat{L} \rightarrow L$$

with a splitting of the same type, in which the coefficient subgroup is elliptic. Since L is a limit group there exists a witness $h : L \rightarrow \Gamma$, and the homomorphism $h \circ \hat{\eta}$ is easily seen to be a witness for the splitting of \hat{L} . By Remark 5.8, if \hat{L} has presentation $\langle X, \mathcal{A} \mid \mathcal{R} \rangle$ then the identity

$$\mathcal{R}(\underline{x}, \underline{a}) \implies \Sigma(\underline{x}, \underline{a})$$

holds in Γ . Therefore, there does exist such a group \hat{L} and so a systematic search will eventually find it and the proof that it has the required properties. \square

Likewise, we have the following theorem.

Theorem 5.18. *Let Γ be a torsion-free hyperbolic group and $\Sigma = 1, \Lambda \neq 1$ a finite, coefficient-free system of equations and inequations. There exists a Turing machine that takes as input Σ and Λ and terminates if and only if the set of conjugacy classes of solutions $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$ is infinite.*

Proof. The proof is very similar to the proof of Theorem 5.17. The Turing machine systematically enumerates finite presentations of the form

$$\hat{L} \cong \langle X \mid \mathcal{R} \rangle.$$

For each presentation, the Turing machine checks to see whether \hat{L} satisfies the following conditions.

- (1) The presentation should exhibit a free or essential cyclic splitting Θ of \hat{L} , or exhibit \hat{L} as an abelian group.
 - (a) In the split case, vertex groups which are not visibly abelian should be seen to be non-abelian via a homomorphism to Γ . Furthermore, there should exist a witness h for Θ that satisfies the inequations $\Lambda(h(\underline{x})) \neq 1$ —again the existence of a witness can be checked by solving a finite system of equations and inequations over Γ .
 - (b) In the abelian case, there should exist a non-trivial element of $\mathcal{O}(\hat{L}, \Gamma; \Lambda)$.

(2) Finally, the sentence

$$\forall \underline{x} \ (\mathcal{R}(\underline{x}) \implies \Sigma(\underline{x}))$$

must hold in Γ .

If such a group \hat{L} is found then $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$ is infinite, by Proposition 5.13 or by Remark 5.14. But, since

$$\mathcal{R}(\underline{x}) \implies \Sigma(\underline{x})$$

it follows that $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$ is infinite. That there exists such a limit group quotient if $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$ is infinite follows exactly as in the proof of Theorem 5.17. \square

6. IF THERE ARE FINITELY MANY

In this section we describe the other halves of the algorithms in the statements of Theorem A and Theorem B. We start with the case when $\text{IHom}_\Gamma(H_\Sigma, \Gamma; \Lambda)$ is finite, as this is much easier.

6.1. Finitely many solutions.

Theorem 6.1. *Let Γ be a torsion-free hyperbolic group. There is an algorithm which takes as input a finite system of equations and inequations (with coefficients) over Γ and terminates if and only if the system has only finitely many solutions.*

If it terminates, it provides a list of the solutions.

Proof. Let $\Sigma = 1$, $\Lambda \neq 1$ be the finite system.

According to Sela [28, Theorem 7.12] (or alternatively [8, Theorem 0.1]), we can determine if there is a solution or not. If there is not, we terminate with the empty list.

If there is a solution, we search for it systematically. Since we know that there is a solution, we will eventually find it. The fact that there is a solution which does not appear on a given finite list is easily encoded into a finite system of equations and inequations, and so we can determine if there is a new solution.

Continuing in this manner, we will clearly discover if there are only finitely many solutions to $\Sigma = 1, \Lambda \neq 1$, and in case we discover this we will find a list of all of the solutions. \square

Combining Theorem 6.9 with Theorem 5.17 completes the proof of Theorem A.

6.2. Finitely many conjugacy classes of solutions. We now turn to the question of recognizing when $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$ is finite. This is substantially more difficult than recognizing when $\text{IHom}_\Gamma(H_\Sigma, \Gamma; \Lambda)$ is finite. In spirit, we want to find solutions that are shortest in their conjugacy classes. This is in general too much to ask, but the results of [9] provide us with a slightly weakened notion that is an adequate substitute. We refer the reader to [9, Section 4] for notation and terminology. Note that we are not concerned with primary peripheral structures in this paper, and we can take them to be empty.

Definition 6.2. *Let H be a group with fixed generating set X . A homomorphism $\psi : H \rightarrow \Gamma$ is compatible if it is injective on the ball of radius 8 in H .*

Compatible homomorphisms are problematic for two reasons. First, a conjugacy class $[\phi] \in \mathcal{O}(H_\Sigma, \Gamma; \Lambda)$ may not have a compatible representative. Secondly, we may not be able to calculate the ball of radius 8, as we have no solution to the word problem in H_Σ . We will circumvent these by supplementing Σ and Λ in order to determine the ball of radius 8.

Definition 6.3. *Let k be a natural number, and let $\mathcal{B}_{F(X)}(k)$ be the ball of radius k in the free group $F(X)$. A system of equations and inequations $\Sigma = 1, \Lambda \neq 1$ forces the ball of radius k if there is a subset $S \subseteq \mathcal{B}_{F(X)}(k)$ such that $S \subseteq \Sigma$ and $\mathcal{B}_{F(X)}(k) \setminus S \subseteq \Lambda$.*

There are finitely many subsets S of the ball of radius 8 in $F(X)$. For such an S , let $\Sigma_S = \Sigma \cup S$ and let $\Lambda_S = \Lambda \cup (\mathcal{B}_{F(X)}(8) \setminus S)$. For any solution ϕ to $\Sigma = 1, \Lambda \neq 1$ there is a unique $S \subseteq \mathcal{B}_{F(X)}(8)$ such that ϕ satisfies $\Sigma_S = 1$ and $\Lambda_S \neq 1$. Therefore we have the following remark.

Remark 6.4. *For $S \subseteq \mathcal{B}_{F(X)}(8)$, let $\eta_S : H_\Sigma \rightarrow H_{\Sigma_S}$ be the obvious epimorphism. The system $\Sigma_S = 1, \Lambda_S \neq 1$ forces the ball of radius 8. Furthermore,*

$$\mathcal{O}(H_\Sigma, \Gamma; \Lambda) = \coprod_{S \subseteq \mathcal{B}_{F(X)}(8)} \eta_S^* \mathcal{O}(H_{\Sigma_S}, \Gamma; \Lambda_S)$$

Moreover, any representative ϕ of a conjugacy class $[\phi] \in \mathcal{O}(H_{\Sigma_S}, \Gamma; \Lambda_S)$ is a compatible homomorphism $H_{\Sigma_S} \rightarrow \Gamma$.

Note that $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$ is finite if and only if $\mathcal{O}(H_{\Sigma_S}, \Gamma; \Lambda_S)$ is finite for each $S \subseteq \mathcal{B}_{F(X)}(8)$.

The property Ω defined in [9, Remark 4.8] provides an adequate means of recognizing short homomorphisms in conjugacy classes. The

property Ω applies to *acceptable lifts* of a homomorphism. See [8] for the definition of an acceptable lift and also [9] for more details.

Definition 6.5. *Fix a pair of elements $a, b \in \Gamma$ that do not commute. As Γ is non-elementary we can always find such a pair. A homomorphism $\psi : H_\Sigma \rightarrow \Gamma$ is fairly short if it is compatible and has an acceptable lift $\tilde{\psi} : \mathcal{B}_{H_\Sigma}(2) \rightarrow F$ that satisfies Ω . The homomorphism ψ is very short if it is compatible and every acceptable lift $\tilde{\psi} : \mathcal{B}_{H_\Sigma}(2) \rightarrow F$ satisfies Ω .*

Remark 6.6. *Here are some important observations.*

- (1) *Note that Ω , as defined in [9, Remark 4.9], depends on the choice of a and b .*
- (2) *The condition Ω is a boolean combination of normalized rational constraints on the lift $\tilde{\psi}$.*

By [9, Lemma 4.7 and Proposition 4.8], we obtain the following.

Lemma 6.7. *Any conjugacy class of compatible homomorphisms $H_\Sigma \rightarrow \Gamma$ contains at least one very short homomorphism, and at most finitely many fairly short homomorphisms.*

Fairly short and very short homomorphism are useful because, as they are characterized by a boolean combination of normalized rational constraints on acceptable lifts, we can follow [8] in exploiting the work of [10] to identify all the very short solutions to a system of equations and inequations. In this terminology, we can deduce the following from [8, Proposition 1.5].

Proposition 6.8. *There exists an algorithm that, given a finite presentation for a torsion-free hyperbolic group Γ and a system of equations $\Sigma = 1$ and inequations $\Lambda \neq 1$ with coefficients in Γ always terminates. If it answers “yes” then there is a fairly short solution to $\Sigma = 1, \Lambda \neq 1$. If it answers “no” then there is no very short solution.*

This is what we need in order to test whether there are only finitely many conjugacy classes of solutions.

Theorem 6.9. *Let Γ be a torsion-free hyperbolic group. There is an algorithm that takes as input a finite system of equations $\Sigma = 1$ and inequations $\Lambda \neq 1$ (without coefficients) and terminates if and only if there are only finitely many conjugacy classes of solutions to the system. In case the algorithm terminates, it provides a list consisting of exactly one representative of each conjugacy class of solutions.*

Proof. In fact, we will describe an algorithm that provides a list containing at least one, and at most finitely many, representatives of each conjugacy class of solutions to $\Sigma = 1, \Lambda \neq 1$. It is easy to pare down such a list so that it contains precisely one representative of each conjugacy class of solutions, using a solution to the simultaneous conjugacy problem in Γ .

For each subset $S \subseteq \mathcal{B}_{F(X)}(8)$ in turn, consider the system $\Sigma_S = 1, \Lambda_S \neq 1$. By Remark 6.4, $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$ is infinite if and only if each of the corresponding $\mathcal{O}(H_{\Sigma_S}, \Gamma; \Lambda_S)$ are infinite. We may therefore assume that $\Sigma = 1, \Lambda \neq 1$ forces the ball of radius 8. In particular, all solutions are compatible.

Let us now describe the procedure for producing the list. Let l be our list of solutions, which initially is empty. Add to $\Sigma = 1, \Lambda \neq 1$ the stipulation that no solution should be equal to a solution on the list l , creating the new system of equations and inequations $\Sigma = 1, \Lambda' \neq 1$. Since we are only adding inequations, we have not changed the notions of very short and fairly short. Now apply the algorithm from Proposition 6.8 to the system $\Sigma = 1, \Lambda' \neq 1$. If the algorithm from Proposition 6.8 says “no” then there is no very short solution to $\Sigma = 1$ and $\Lambda' \neq 1$, and the procedure terminates with the list l . If the algorithm says “yes” then there is a fairly short solution. A naive search will eventually find such a solution. It cannot be on the list l , so add it to l and repeat the process.

By Lemma 6.7, every conjugacy class of solutions has a very short representative. It follows that, if $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$ is infinite then the algorithm will not terminate. On the other hand Lemma 6.7 states that every conjugacy class contains only finitely many fairly short representatives, so if $\mathcal{O}(H_\Sigma, \Gamma; \Lambda)$ is finite then the algorithm will eventually terminate and output the list l . \square

Combining Theorem 6.9 with Theorem 5.18 completes the proof of Theorem B.

7. IMMUTABLE SUBGROUPS

7.1. Enumerating immutable subgroups.

Definition 7.1. *Let G be a group. A finitely generated subgroup H of G is called immutable if there are finitely many injections $\phi_1, \dots, \phi_N : H \rightarrow G$ so that any injection $\phi : H \rightarrow G$ is conjugate to one of the ϕ_i .*

We will see that all Γ -limit groups are either strict or immutable subgroups of Γ . Therefore, in order to understand Γ -limit groups, or even just subgroups of Γ , we need to study immutable subgroups.

This is one of the big differences between Γ -limit groups where Γ is a torsion-free hyperbolic group and ordinary limit groups defined over a free group. Clearly, there are no non-trivial immutable subgroups of a free group.

Lemma 7.2. *A finitely generated subgroup of Γ is immutable if and only if it does not admit a non-trivial free splitting or an essential splitting over \mathbb{Z} .*

Proof. Let H be a finitely generated subgroup of Γ . If $H \cong \mathbb{Z}$ then H is not immutable, since \mathbb{Z} splits over 1.

If H admits a non-trivial free or essential cyclic splitting then, taking the inclusion $h : H \hookrightarrow \Gamma$ as a witness, it follows immediately by Proposition 5.13 that there are infinitely many conjugacy classes of homomorphisms $H \rightarrow \Gamma$. As each is obtained by pre-composing h with an automorphism they are all injective, and so H is not immutable.

Suppose, then, that H admits infinitely many conjugacy classes of embeddings into Γ . We can therefore choose a sequence of pairwise non-conjugate embeddings $\{\phi_i : H \rightarrow \Gamma\}$. After passing to a subsequence, the sequence ϕ_i can be taken to be stable. As $H = H/\overline{\text{Ker}}(\phi_i)$, this realizes H as a *strict* Γ -limit group. It now follows from Theorem 3.9 that H splits essentially over a cyclic subgroup. \square

Remark 7.3. *The proof of Lemma 7.2 shows that any Γ -limit group is either strict or an immutable subgroup of Γ , and the two cases are mutually exclusive.*

Example 7.4. *Let Γ be the fundamental group of a closed hyperbolic n -manifold where $n \geq 3$. Then Γ admits no non-trivial abelian splitting (see, for example, [1, Theorem 1.6(i)]) and so, by Lemma 7.2, is an immutable subgroup of itself.*

In [19], a version of the Rips construction is given in which the kernel has Kazhdan's property (T), and in particular is immutable. This provides many further examples of finitely generated (but not finitely presentable) immutable subgroups.

The following lemma provides a slightly weaker characterization of immutable subgroups. The advantage is that it can be recognized algorithmically.

Lemma 7.5. *A finitely generated subgroup H of Γ is immutable if and only if there are finitely many homomorphisms $\rho_1, \dots, \rho_K : H \rightarrow \Gamma$ and an integer $D \geq 1$, so that for any homomorphism $\rho : H \rightarrow \Gamma$, either*

- (1) ρ is conjugate to one of the ρ_i ; or
- (2) ρ is not injective on the ball of radius D about 1 in H .

Proof. If H is not immutable then it admits infinitely many conjugacy classes of embeddings, and so clearly does not satisfy the dichotomy in the statement of the lemma.

On the other hand, suppose that H does not satisfy the statement of the lemma. Then there is a sequence $\{\phi_i : H \rightarrow \Gamma\}$ of non-conjugate homomorphisms so that ϕ_i is injective on the ball of radius i about 1 in H . We may now pass to a stable subsequence with trivial stable kernel and deduce that H is a strict Γ -limit group and therefore not immutable. \square

Note that we do not insist that the homomorphisms ρ_i are injective. However, since H is assumed to be a subgroup of Γ , necessarily at least one of the ρ_i is an embedding.

Example 7.6. *Fix a generating set X for H an immutable subgroup of Γ (e.g. Example 7.4). The property that a map from X into Γ extends to a homomorphism $h : H \rightarrow \Gamma$ that is injective on the ball of radius D in the Cayley graph of H with respect to X is easily encoded as a system of (possibly infinitely many) equations and finitely many inequations over Γ . By Lemma 7.5, there are only finitely many conjugacy classes of solutions. By Theorem 4.4, we have an example of a finite system of equations and inequations with only finitely many conjugacy classes of solutions in Γ .*

Note that if H happens to be finitely presentable (as in Example 7.4), it is not necessary to appeal to Theorem 4.4.

One of the problems with finitely generated subgroups of torsion-free hyperbolic groups is that they need not be finitely presentable. There are also torsion-free hyperbolic groups for which it is undecidable which finite subsets generate finitely presentable subgroups (see [7, Theorem 1]). We do not know the answer to the following question.

Question 7.7. *Let Γ be a torsion-free hyperbolic group. Is there an algorithm that takes as input finite subsets \mathcal{A} of Γ so that $\langle \mathcal{A} \rangle$ is finitely presented and finds a finite presentation for $\langle \mathcal{A} \rangle$?*

We now have the tools to prove that we can recognize algorithmically when a finite set generates an immutable subgroup (although not when it does not!). Theorem C is an immediate corollary.

Proposition 7.8. *There exists a Turing machine that takes as input a finite subset $\mathcal{A} \subseteq \Gamma$ and terminates if and only if the subgroup $H = \langle \mathcal{A} \rangle$ is immutable.*

Proof. Our approach is to construct a sequence of epimorphisms $\hat{H}_D \rightarrow H$ (for any $D \in \mathbb{N}$) and to test the criterion of Lemma 7.5 on each \hat{H}_D .

Let \hat{H}_D be the group generated by S with relations equal to all the loops in the ball of radius D in H . So there is an epimorphism $\eta_D : \hat{H}_D \rightarrow H$ that is an isometry between the balls of radius D . Homomorphisms $\hat{H}_D \rightarrow \Gamma$ that are injective on the ball of radius D in \hat{H}_D can be characterized precisely by a system of equations and inequations over Γ —the equations are the relations of \hat{H}_D and the inequations separate the points in the ball of radius D . (Although we do not have a solution to the word problem in \hat{H}_D , we can still construct its ball of radius D since it is, by definition, isometric to the ball of radius D in H .) Using the algorithm of Theorem B we can now determine whether this system of equations and inequations has finitely many conjugacy classes of solutions. If it does then it follows immediately that H is immutable.

Checking this for each D in turn, we have a process that, if it terminates, confirms that H is immutable. It remains to show that the process terminates whenever H is immutable. By Theorem 4.5, $\text{Hom}(\hat{H}_D, \Gamma) = \text{Hom}(H, \Gamma)$ for all sufficiently large D —that is, in our terminology, \hat{H}_D is a Γ -approximation for H . For large enough D we will also have that D satisfies the conclusions of Lemma 7.5. For such a D , there are only finitely many conjugacy classes of homomorphisms $\hat{H}_D \rightarrow \Gamma$ that are injective on the ball of radius D . Therefore, if H is immutable then the process will terminate. \square

7.2. Questions. We have proved that the set of immutable subgroups of Γ is recursively enumerable.

Question 7.9. *Is the set of immutable subgroups of Γ recursive?*

Here a class of groups is *recursive* if there is an algorithm that enumerates the class of groups so that the list contains a single representative of each isomorphism class. This is essentially the same as asking whether the isomorphism problem is solvable for immutable subgroups. One would expect the answer to be negative, as the isomorphism problem for finitely generated subgroups of hyperbolic groups is unsolvable.

Applying the construction of [19] to a group with unsolvable word problem, we obtain the following.

Proposition 7.10. *There are torsion-free hyperbolic groups Γ for which the membership problem for finitely generated immutable subgroups is unsolvable.*

However, no immutable subgroup constructed in this way can be finitely presentable.

Question 7.11. *Is the membership problem for finitely presentable immutable subgroups solvable?*

Question 7.11 appears to be related to the isomorphism problem for finitely presented subgroups of torsion-free hyperbolic groups.

Question 7.12. *Let Γ be a (torsion-free) hyperbolic group. Is there an algorithm that takes as input a finite subset of Γ and decides whether or not the subgroup generated is immutable?*

Decision problems about finitely generated subgroups of hyperbolic groups are rarely solvable (see, for instance, [23], [7]). Nevertheless, the results of this paper reduce Question 7.12 to determining if a subgroup of Γ splits freely or essentially over \mathbb{Z} .

REFERENCES

- [1] I. Belegradek, Aspherical manifolds with relatively hyperbolic fundamental groups, *IJAC* **18** (2008), 97–110.
- [2] M. Bestvina, Degenerations of the hyperbolic space, *Duke Math. J.* **56** (1988), 143–161.
- [3] M. Bestvina, \mathbb{R} -trees in topology, geometry, and group theory, *Handbook of geometric topology*, 55–91, North Holland, 2002.
- [4] M. Bestvina and M. Feighn, Notes on Sela’s work: Limit groups and Makanin-Razborov diagrams, preprint.
- [5] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1999.
- [6] M.R. Bridson and G.A. Swarup, On Hausdorff-Gromov convergence and a theorem of Paulin, *Enseign. Math.* **40** (1994), 267–289.
- [7] M. Bridson and D. Wise, Malnormality is undecidable in hyperbolic groups, *Israel J. Math.* **124** (2001), 313–316.
- [8] F. Dahmani, Existential questions in (relatively) hyperbolic groups, *Israel J. Math.*, to appear.
- [9] F. Dahmani and D. Groves, The Isomorphism Problem for toral relatively hyperbolic groups, *Publ. Math. IHES* **107** (2008), 211–290..
- [10] V. Diekert, C. Gutiérrez, C. Hagenah, The existential theory of equations with rational constraints in free groups is PSPACE-complete. STACS 2001 (Dresden), 170–182, Lecture Notes in Comput. Sci., 2010, Springer, 2001.
- [11] D. Groves, Limit groups for relatively hyperbolic groups, II: Makanin-Razborov diagrams, *Geom. and Topol.* **9** (2005), 2319–2358.
- [12] D. Groves and H. Wilton, Enumerating limit groups, *Groups, Geom. and Dynamics* **3** (2009), 389–399.
- [13] D. Groves and H. Wilton, Limit groups over hyperbolic groups, in preparation.
- [14] Vincent Guirardel. Actions of finitely generated groups on \mathbb{R} -trees. *Ann. Inst. Fourier*, **58** (2008), 159–211.
- [15] O. Kharlampovich and A. Miasnikov, Irreducible affine varieties over a free group II, *Journal of Algebra* **200** (1998), 517–570.

- [16] O. Kharlampovich and A. Miasnikov, Elementary theory of nonabelian free groups, *J. Algebra* **302** (2006), 451–552.
- [17] G.S. Makanin, Equations in a free group. (Russian) *Izv. Akad. Nauk SSSR Ser. Nat.* **46** (1982), 1199–1273; translation in *Math. USSR-Izv.* **21** (1983), 546–582.
- [18] G.S. Makanin, Decidability of the universal and positive theories of a free group. (Russian) *Izv. Akad. Nauk SSSR Ser. Mat.* **48** (1984), 735–749; translation in *Math. USSR-Izv.* **25** (1985), 75–88.
- [19] Y. Ollivier and D. Wise, Kazhdan groups with infinite outer automorphism group, *Trans. AMS* **360** (2007), 1959–1976.
- [20] Frédéric Paulin. *Outer automorphisms of hyperbolic groups and small actions on \mathbb{R} -trees*, volume 19 of *Math. Sci. Res. Inst. Publ.*, pages 331–343. Springer, New York, 1991.
- [21] A. Razborov, On systems of equations in a free group, *Math. USSR Izvestiya* **25** (1985), 115–162.
- [22] A. Razborov, On systems of equations in a free group, Ph.D. thesis, *Steklov Math. Institute* (1987).
- [23] E. Rips, Subgroups of small cancellation groups, *Bull LMS* **14** (1982), 45–47.
- [24] E. Rips and Z. Sela, Canonical representatives and equations in hyperbolic groups. *Invent. Math.* **120** (1995), no. 3, 489–512.
- [25] Z. Sela. Acylindrical accessibility for groups. *Invent. Math.*, **129** (1997) 527–565.
- [26] Z. Sela, Diophantine geometry over groups, I: Makanin-Razborov diagrams, *Publ. Math. IHES* **93** (2001), 31–105.
- [27] Z. Sela, Diophantine geometry over groups. VI. The elementary theory of a free group, *GAF* **16** (2006), 707–730.
- [28] Z. Sela, Diophantine geometry over groups, VII: The elementary theory of hyperbolic groups, *Proc. LMS* **99** (2009), 217–273..

DANIEL GROVES, MSCS UIC 322 SEO, M/C 249, 851 S. MORGAN ST.,
CHICAGO, IL 60607-7045, USA

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

HENRY WILTON, DEPT OF MATHEMATICS, 1 UNIVERSITY STATION C1200,
AUSTIN, TX 78712, USA

E-mail address: henry.wilton@math.utexas.edu