AMALGAMATED FREE PRODUCTS, HNN EXTENSIONS, AND DECISION PROBLEMS

ESME BAJO

Abstract. This paper provides an introduction to combinatorial group theory, culminating in an exploration of amalgamated free products, HNN extensions, and their actions on trees. This approach to amalgamated free products and HNN extensions will be in the context of Bass-Serre Theory and fundamental groups of graphs of groups, and we will then apply these abstract notions to a wider range of results, namely the Novikov-Boone Theorem on the undecidability of the word problem and the Adian-Rabin Theorem on the undecidability of Markov properties.

Contents

1. Introduction 1
2. Some Basic Definitions 2
3. Free Groups Actions and Free Groups 4
4. Fundamental Groups of Graphs 6
5. Amalgamated Free Products and HNN Extensions 8
6. Fundamental Groups of Graphs of Groups 11
7. The Novikov-Boone Theorem 14
8. The Adian-Rabin Theorem 16
Acknowledgments 18
References 18

1. Introduction

This paper provides proofs of some important theorems in computability theory, namely the Novikov-Boone and Adian-Rabin theorems:

Theorem 1.1 (Novikov-Boone)**.** *There exists a finitely presented group $G = \langle X \mid R \rangle$ for which the word problem in G is undecidable.*

Theorem 1.2 (Adian-Rabin)**.** *There is no algorithm to decide whether or not a finitely presented group satisfies a given Markov property.*

We will use Theorem 1.1 to prove Theorem 1.2. We will formally define a Markov property later, but for now note that being trivial, being finite, being free, being cyclic, being simple, being solvable, and being torsion-free are all Markov properties of finitely presented groups, so the Adian-Rabin Theorem is a powerful theorem.

While these theorems delve into computability theory, we will provide proofs based on two concepts constructed via combinatorial group theory: amalgamated free products and HNN extensions. We build up to these constructions with an

introduction to combinatorial group theory, including definitions of graphs, group actions on graphs, and fundamental groups of graphs, and an exploration of Bass-Serre theory.

Jean-Pierre Serre developed the notion of graphs of groups to describe certain algebraic groups. He related group actions on trees to iterated applications of amalgamated free products and HNN extensions, and defined the *fundamental group of a graph of groups*, generalizing these two concepts. We will provide rigorous definitions of these notions, and our discussion will culminate in the following fundamental theorem:

Theorem 1.3 (Structure Theorem)**.** *If $G = \pi_1(\mathbb{G}, Y, T)$, then G acts without inversion of edges on a tree X such that $Y \cong G \setminus X$ and the stabilizers of the vertices and edges of the tree X are conjugate to the images (under the canonical morphism) of G_v, $v \in Y^0$, and $\alpha_e(G_e)$, $e \in Y^1$, respectively.*

Conversely, if G acts on a tree X, then G is isomorphic to the fundamental group $\pi_1(\mathbb{G}, G \setminus X, T)$, where T is a maximal subtree of $G \setminus X$ and the vertex and edge groups correspond to the stabilizers of the vertices and edges of X, respectively.

We build to an application of Bass-Serre theory to computability theory, originally discovered by Max Dehn. Historically, fundamental groups have been used in many areas of mathematics, including knot theory in topology. In the early 1900s, Dehn proposed the problem of deciding whether or not two knots are the same. He embedded knots in 3-space and considered the fundamental groups of the complements of these knots (which are not isomorphic for distinct knots). He discovered that these knot theory problems could be generalized to problems about finitely presented groups, and identified three problems: the word problem, the conjugacy problem, and the isomorphism problem [3].

In this paper, we provide a proof of the undecidability of the word problem for finitely presented groups (the Novikov-Boone theorem) reliant on an amalgamated free product, and a subsequent proof of the Adian-Rabin theorem based on successive amalgamated products and HNN extensions. This theorem yields the undecidability of the isomorphism problem as an easy corollary. These proofs rely on our group theoretic definitions rather than an intensive treatment of computability theory. Therefore, we work off of some assumptions, notably the notion of undecidability and the existence of a recursively enumerable set that is not recursive.

2. Some Basic Definitions

First, we construct some preliminary notions of combinatorial group theory, beginning with Serre's definition of a graph and a corresponding description of a group acting on a graph. We assume prior knowledge of general group actions.

Definition 2.1. A *graph* X is a (nonempty) set of vertices X^0, a set of edges X^1, and the following three maps:

i) $\alpha : X^1 \to X^0$ sending an edge to its initial vertex

ii) $\omega : X^1 \to X^0$ sending an edge to its final vertex

iii) $^- : X^1 \to X^1$ sending each edge $e \in X^1$ to its inverse $\bar{e}$, where we require $\alpha(e) = \omega(\bar{e})$, $\bar{\bar{e}} = e$, and $\bar{e} \neq e$ for all $e \in X^1$

An *oriented* graph consists of the vertices X^0 and for each $e \in X^1$, a choice of exactly one of $\{e, \bar{e}\}$. We call an edge *positively oriented* if it is in a given orientation and *negatively oriented* if it is not. We denote the set of all positively oriented edges

by X^1_+ and the set of all negatively oriented edges by X^1_-, and we often refer to an orientation by this set X^1_+.

We will be interested in some specific types of graphs. In particular, we define a *connected graph* to be a graph X such that for any distinct $u, v \in X^0$, there exists a path $e_1e_2...e_k$ with $\alpha(e_1) = u$, $\omega(e_k) = v$, and $\omega(e_i) = \alpha(e_{i+1})$ for $1 \leqslant i \leqslant k-1$. We call a path p reduced if there are no subsequences $e\bar{e}$ in p, and we define a *tree* to be a connected graph X such that for any $v \in X^0$, there is no reduced path (of length greater than 0) from v to itself. Note that in a tree, there is a unique reduced path between any two vertices (otherwise, two distinct paths would create a closed path). We will be interested in groups acting on graphs, particularly on trees.

Definition 2.2. A (left) group action of G on a graph X is a pair of (left) actions, $G \curvearrowright X^0$ and $G \curvearrowright X^1$, such that $g\alpha(e) = \alpha(ge)$ and $g\bar{e} = \overline{ge}$ for all $g \in G$ and $e \in X^1$. (It follows that $g\omega(e) = \omega(ge)$ for all $g \in G$ and $e \in X^1$.)

For the purposes of this paper, we require that G act on X *without inversion of edges*, i.e. that for any $g \in G$ and $e \in X^1$, $ge \neq \bar{e}$. But we show that this requirement is not restrictive, because if G acts on X then we can construct the barycentric subdivision $B(X)$ of X so that G acts on $B(X)$ without inversion of edges:

Construction 2.3. Form the barycentric subdivision of a graph X as follows:

i) For each $e \in X^1$, where $\alpha(e) = u$ and $\omega(e) = v$, add a vertex w and replace e with edges e_1 and e_2, where $\alpha(e_1) = u$, $\omega(e_1) = \alpha(e_2) = w$, and $\omega(e_2) = v$.

ii) Set $\overline{e_1} = (\bar{e})_2$ and $\overline{e_2} = (\bar{e})_1$.

iii) For $G \curvearrowright X$, define the action of G on $B(X)$ so that $g(e_1) = (ge)_1$ and $g(e_2) = (ge)_2$.

$$X: \quad u \xrightarrow{\ e\ } v$$

$$B(X): \quad u \xrightarrow{\ e_1\ } w \xrightarrow{\ e_2\ } v$$

Note that if we have a group G acting on X *with* inversion of edges, then we have $g(e) = \bar{e}$ for some $e \in X^1$. But this implies $g(e_1) = (ge)_1 = (\bar{e})_1 = \overline{e_2} \neq \overline{e_1}$ and $g(e_2) = (ge)_2 = (\bar{e})_2 = \overline{e_1} \neq \overline{e_2}$. Hence, we have $G \curvearrowright B(X)$ *without* inversion of edges. $G \curvearrowright B(X)$ behaves (in some sense) like $G \curvearrowright X$.

If a group G acts on a graph X, then for each $x \in X^0 \cup X^1$ we can consider its *orbit* $\mathcal{O}(x) = \{y \in X^0 \cup X^1 \mid gx = y \text{ for some } g \in G\}$. We can now define another type of graph (corresponding to a group action $G \curvearrowright X$):

Definition 2.4. The *factor graph* $G \setminus X$ is the graph with vertices $\mathcal{O}(v)$ for $v \in X^0$ and edges $\mathcal{O}(e)$ for $e \in X^1$, connected such that:

i) $\mathcal{O}(v) = \alpha(\mathcal{O}(e))$ if $gv = \alpha(e)$ for some $g \in G$

ii) $\mathcal{O}(\bar{e})$ is the inverse of $\mathcal{O}(e)$ for all e

(It follows that $\mathcal{O}(v) = \omega(\mathcal{O}(e))$ if $gv = \omega(e)$ for some $g \in G$.)

Example 2.5. Suppose C_4 acts on the graph X (shown on the next page), where 1 acts via a 90° rotation clockwise. There is a single orbit of vertices $\mathcal{O}_v = \mathcal{O}(A) =$

$\{A, B, C, D\}$ and a single orbit of edges $\mathcal{O}_e = \mathcal{O}(e_1) = \{e_1, e_2, e_3, e_4\}$, so the factor graph $C_4 \setminus X$ is a loop:

X :

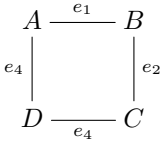

$C_4 \setminus X$: 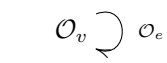

3. Free Groups Actions and Free Groups

In this section, we define free group actions and then reconcile this definition with that of free groups.

Definition 3.1. A group action $G \curvearrowright X$ is *free* if $gv = v$ for any $v \in X^0 \implies g = e$. For example, C_4 acts freely on X in Example 2.5.
(Note that this condition is stronger than an action being faithful. Under a free action a nontrivial element of G cannot fix *any* vertex, whereas under a faithful action a nontrivial element of G cannot fix *all* vertices.)

Suppose we have a group G acting freely (and without inversion of edges) on a tree X. We want to describe the group G. Consider the canonical projection ψ of X onto the factor graph $G \setminus X$. Orient X and $G \setminus X$ such that an edge $e \in X^1$ is positively oriented if and only if its image $\psi(e)$ is positively oriented in $(G \setminus X)^1$. We can choose a maximal subtree T (i.e. a tree containing each vertex of the graph) in $G \setminus X$. Note that every vertex of X is the preimage (under ψ) of some vertex of $G \setminus X$, and thus of some vertex in T. Furthermore, since T is a tree there is a unique path between any two $v, w \in T^0$, and thus once we fix some preimage $v_0 \in X^0$ of a given vertex $v \in T^0$, there is a unique preimage of any other vertex $w \in T^0$. Therefore, we can fix a preimage of some $v \in T^0$ and then lift all of T to a tree T_0 within X, and the elements of G will move T_0 around to the other lifts of T in X.

Now, we consider the set of positively oriented edges in $G \setminus X$ that are *not* in our maximal subtree T. Call this set E. Each $e \in E$ will have an initial vertex in T, so there will be a lift of e to an edge $e_0 \in X^1$ with initial vertex in our fixed T_0 and final vertex outside of T_0. Suppose there were multiple possible final vertices, and we had lifts e_0 and e_0'. Then the action of G taking e_0 to e_0' would fix the initial vertex, and $G \curvearrowright X$ would not be free. Hence, our lift e_0 is unique. Denote the set of such e_0 (positively oriented edges of X with initial vertices in T_0 and final vertices outside of T_0) by E_0.

Any $e_0 \in E_0$ must have a terminal vertex outside of T_0 and since G acts freely on X, we have a unique element $g_{e_0} \neq 1$ such that $\omega(e_0)$ is in $g_{e_0} T_0$. Let S_0 be the set of such elements g_{e_0}, where $e_0 \in E^0$. We will show that the group G satisfies the following definition (with respect to basis S_0):

Definition 3.2. If G is a group and $S \subseteq G$ is linearly ordered such that $S \cap S^{-1} = \varnothing$, then G is a *free group* (with basis S) if each nontrivial $g \in G$ can be uniquely expressed as $g = s_1 ... s_n$ for $s_i \in S \cup S^{-1}$, where $s_i s_{i+1} \neq 1$ for any i.

We have seen that the subtrees gT_0, $g \in G$, in our graph X are disjoint and the set of all such subtrees covers all vertices in X. Since the gT_0 are disjoint, we have a bijection between the elements $g \in G$ and the subtrees gT_0 in X. Therefore, we can consider the graph $\tilde{X}$, where each gT_0 in X is contracted to single vertex in $\tilde{X}$, denoted $\tilde{g}$. Since X is a tree, $\tilde{X}$ is a tree. We now want to define an isomorphism between $\tilde{X}$ and the *Cayley graph* Γ of G and S_0:

Definition 3.3. Construct the Cayley graph Γ of G and S as follows:

Let there be one vertex for every element in G. For all $g \in G$ and $s \in S$, draw an edge e with $\alpha(e) = g$ and $\omega(e) = gs$ (and a corresponding inverse edge $\bar{e}$ with $\alpha(\bar{e}) = gs$ and $\omega(\bar{e}) = g$).

Example 3.4. For example, the image below shows the Cayley graph of the free group generated by two elements a and b (with respect to that basis $\langle a, b \rangle$) [2].

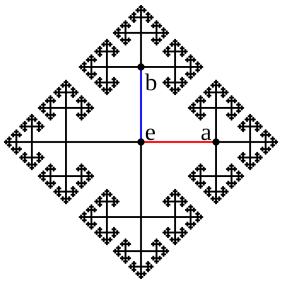

Clearly the map $\tilde{g} \mapsto g$ defines a bijection between the vertices of $\tilde{X}$ and Γ. Now we want to define a bijection between the edges of $\tilde{X}$ and Γ. For any edge in $\tilde{X}$, e cannot have been contained in any subtree gT_0 in X, and thus must have its initial and terminal vertices in distinct copies of T_0. Let $e \in \tilde{X}^1$ be an edge in $\tilde{X}$, where $\alpha(e) = \tilde{g_1}$ and $\omega(e) = \tilde{g_2}$, $g_1 \neq g_2$. Then let $s = g_1 g_2^{-1}$. Map $e \in \tilde{X}^1$ to the edge $e' \in \Gamma$, where $\alpha(e') = g_1$ and $\omega(e') = g_1 s = g_2$. We want to show that this is in fact a bijection. Suppose there is an edge from g to gs, $s \in S_0$, in Γ. Since $s \in S_0$, there is an edge from T_0 to gT_0 in X, and thus an edge from $\tilde{1}$ to $\tilde{g}$ in $\tilde{X}$. Clearly, we have constructed an isomorphism between $\tilde{X}$ and Γ. To complete our proof that G is free with respect to basis S_0, we show that the Cayley graph of G and S_0 is a tree only if G is free with respect to basis S_0.

If Γ is a tree, then for any $g \in G$, $g \neq 1$, there is a unique reduced path from 1 to g in Γ. Let this be given by $e_1 ... e_n$, $n \geqslant 1$. Put $\alpha(e_i) = h_i$ and $\omega(e_i) = h_i s_i$ for each i, where all $h_i \in G$ and all $s_i \in S \cup S^{-1}$. Then we can write

$$g = \omega(e_n) = h_n s_n = \alpha(e_n) s_n = \omega(e_{n-1}) s_n = h_{n-1} s_{n-1} s_n = ...$$

$$= \alpha(e_1) s_1 ... s_n = s_1 ... s_n$$

Now we just need to show that this expression is unique. Suppose not. Then we have two distinct paths from 1 to g, contradicting Γ being a tree.

To recap, we have shown that if we have a free action $G \curvearrowright X$, where X is a tree, then we can lift a maximal subtree T in $G \backslash X$ to a subtree T_0 in X. This subtree is translated to other copies of T_0 by elements of G. Therefore, we can contract these gT_0 to single vertices, creating a new tree $\tilde{X}$, which is isomorphic to the Cayley graph Γ of G with respect to the basis S_0 (where the elements of S_0

were the elements taking vertices in T_0 to adjacent vertices outside of T_0). Finally, we showed that Γ being a tree must imply that G is free with respect to that basis.

This shows that a group G that acts freely (and without inversion of edges) on a tree is free. We will now show that these two notions are actually equivalent:

Proposition 3.5. *A group G acts freely (and without inversion of edges) on a tree if and only if it is free.*

Proof. We have already shown one direction, so we are left to show that all free groups act freely on a tree. Suppose G is a free group with basis S. Consider the Cayley graph Γ of G and S. Since S generates G, Γ is connected. We claim that Γ is in fact a tree. If not, then there exists a closed path $e_0...e_n$ with $\alpha(e_0) = \omega(e_n) = g$ for some $g \in G$. But then $g = gs_0^{\varepsilon_0}...s_n^{\varepsilon_n}$, where each $s_i \in S$ and $\varepsilon_i \in \{-1, 1\}$. But this contradicts G being free with respect to S. Hence Γ is a tree. Now we claim that G acts (via left multiplication) freely and without inversion of edges on the tree Γ. Suppose some $g \in G$ fixes a vertex $v_h \in \Gamma^o$ (where v_h corresponds to element $h \in G$). This implies $gh = h \implies g = 1$. So the action must be free. Furthermore, this action is without inversion of edges. Otherwise, we would have some $g_2 = g_1 s$ and $g_1 = g_2 s$ for $s \in S$. Then $g_1 = g_1 s^2 \implies s^2 = 1$, which contradicts the freeness of G with respect to S. □

This result yields the following corollary:

Corollary 3.6 (The Nielsen-Schreier Theorem)**.** *Any subgroup of a free group is free.*

Proof. Suppose H is a subgroup of the free group G (with basis S). We have shown that G acts freely (and without inversion of edges) on the Cayley graph Γ of G and S. But $H \leqslant G$, so clearly H does as well. Thus, H acts freely on the tree Γ, and must itself be free by Proposition 3.5. □

4. Fundamental Groups of Graphs

So far, we have studied group actions on trees, but we now want to look at more general group actions on graphs. We examine graphs with nondegenerate closed paths using the notion of fundamental groups, constructed below:

Given a connected graph X and a fixed vertex $x \in X^0$, consider the set P of all paths $e_0...e_n$ in X such that $\alpha(e_0) = \omega(e_n) = x$. We define the following equivalence relation on P: any two paths $e_0...e_{i-1}e_i\bar{e_i}e_{i+1}...e_n$ and $e_0...e_{i-1}e_{i+1}...e_n$ are equivalent. Denote the equivalence class of a path p by $[p]$. We note that the set of *reduced* paths (where we don't go back and forth along an edge) is a set of representatives of these equivalence classes. To make the set of equivalence classes of P a group, we define multiplication as follows: for two paths $p = e_0...e_n$ and $q = f_0...f_m$ from x to x, let $pq = e_0...e_nf_0...f_m$ and $[p][q] = [pq]$. The identity element is the equivalence class containing the degenerate path of length 0 from x to x.

Definition 4.1. This group of equivalence classes of P is the *fundamental group* $\pi_1(X, x)$ of the graph X with respect to the vertex x.

This notion is a way of describing the closed paths in a graph. For example, in a tree, there is a unique reduced path between any two vertices, so there are no non-degenerate reduced paths from a vertex to itself. Hence, there is only one

equivalence class of paths from a vertex to itself. The fundamental group of a tree with respect to any of its vertices is thus the trivial group.

In this way, subtrees in a graph do not contribute to its fundamental group, as traversing them and returning to the original vertex requires going back and forth along an edge. But this will be in the same equivalence class as the same path that does not traverse that edge in the tree. This suggests the following proposition:

Proposition 4.2. *Let X be a connected graph, X^1_+ be an orientation of X, $x \in X^0$ be any vertex, and T be a maximal subtree of X. For each $v \in X^0$ we know that there is a unique path from x to v in T, which we denote p_v. For each $e \in X^1$, let $p_e = p_{\alpha(e)} e p^{-1}_{\omega(e)}$. Then $\pi_1(X, x)$ is a free group with basis $S = \{[p_e] \mid e \in X^1_+ - T^1\}$.*

Proof. First, we want to show that any element of the fundamental group $\pi_1(X, x)$ can be written as the product of elements of S. Suppose $[p] \in \pi_1(X, x)$, where $[p]$ is the equivalence class of a closed path $p = e_1 ... e_n$ from x to x in X. Note that for all i, we have $\omega(e_i) = \alpha(e_{i+1})$, so $p_{\omega(e_i)} = p_{\alpha(e_{i+1})} \implies p^{-1}_{\omega(e_i)} p_{\alpha(e_{i+1})} = 1$, and thus:

$$[p] = [p_{e_1} p_{e_2} ... p_{e_n}] = [p_{e_1}][p_{e_2}]...[p_{e_n}]$$

If $[p_{e_i}] \notin S$, then e_i is contained in the subtree T, and thus, $[p_{e_i}] = [1]$. Thus, any $[p]$ can be written as the product of elements of S.

To show that $\pi_1(X, x)$ is free with respect to S, we now just need to show that this representation is unique. Suppose we can write

$$[p] = [p_{e_1} ... p_{e_n}] = [p_{e_1}]...[p_{e_n}],$$

where each $[p_{e_i}] \in \pi_1(X, x)$ (or equivalently, each $e_i \in X^1_+ - T^1$). Note that for each $i \in \{1, ..., n\}$, the path p_{e_i} is fully contained in T, except at the edge e_i. Thus it does not contain any edge e_j, $e_i \neq e_j$, for $j \in \{1, ...n\}$. Therefore, if we reduce the path $p_{e_1} ... p_{e_n}$, there are no reductions along any e_i, and we get a reduced path of the form $t_0 e_1 t_1 ... t_{n-1} e_n t_n$, where each t_k is a reduced path in T. But each equivalence class in $\pi_1(X, x)$ contains exactly one reduced path, so the $e_1, ..., e_n$ are *uniquely* determined by p. □

Example 4.3. Consider the following graph X, and the associated orientation and selection of a maximal subtree T (bolded) shown below it:

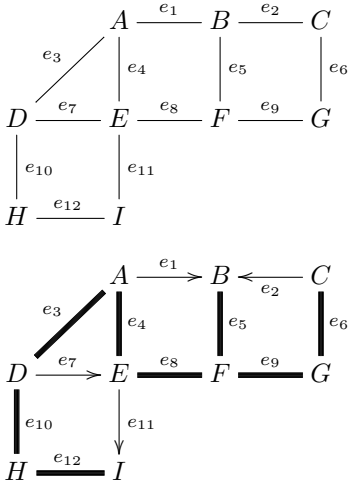

By Proposition 4.2, $\pi_1(X, A) \cong F_4$. In order to think about the fundamental group of graph X, imagine laying a piece of string along a path in X and pulling it taut. As long as we stay within the maximal subtree, we can pull the string back to point A, but if we cross e_1, e_2, e_7, or e_{11} and pull it taut, the string will catch around a corresponding closed path.

We have shown that given connected graph X and a maximal subtree T, $\pi_1(X, x)$ (for $x \in X^0$) is freely generated by the paths from x to x that cross exactly one positively oriented edge outside of T. But note that in Section 3, we showed that if a group G acts freely on a a tree X, then G is free and generated by a set of cardinality equal to the cardinality of the set of positively oriented edges of the factor graph $G \setminus X$ laying outside a fixed maximal tree. We have thus shown the following:

Theorem 4.4 (Fundamental Theorem of Bass-Serre Theory)**.** *Let G be a group acting freely and without inversion of edges on a tree X. Fix any vertex $v \in X^0$. Then G is isomorphic to the fundamental group $\pi_1(G \setminus X, v)$.*

Serre generalized this theorem to an analogous theorem for *fundamental groups of graphs of groups* (Section 6). But before we do this, we construct amalgamated free products and HNN extensions, which we will later recover as examples of this generalized notion.

5. Amalgamated Free Products and HNN Extensions

We must first define the *free product* of two groups A and B, which is a group of *normal forms*. Given groups A and B such that $A \cap B = 1$ (note that this is not a restriction as we may take isomorphic copies of any A and B with nontrivial intersection), a normal form is an expression $x_1...x_n$ such that for all i, $1 \leqslant i \leqslant n$, we have that $x_i \in A \cup B$ and any adjacent x_i, x_{i+1} are not both in A and not both in B. n is the length of such a normal form and the identity element is said to have length 0. In order to make the set of normal forms into a group, we define multiplication on the set of all normal forms as follows:
Let $x = x_1..., x_n$ and $y = y_1...y_m$ be normal forms of A and B. Put $x \cdot y =$

i) x if $y = 1$

ii) y if $x = 1$

iii) $x_1...x_n y_1...y_m$ if x_n and y_1 are in different groups

iv) $x_1...x_{n-1} z y_2...y_m$, where $z = x_n y_1$, if x_n and y_1 are in the same group and $x_n y_1 \neq 1$

v) $(x_1...x_{n-1}) \cdot (y_2...y_m)$ if x_n and y_1 are in the same group and $x_n y_1 = 1$

The set of all normal forms is clearly a group under this multiplication, and this group is called the *free product* of A and B, and is denoted by $A * B$.

Example 5.1. Consider the free product $\mathbb{Z} * \mathbb{Z}$. Let the first copy of $\mathbb{Z}$ be generated by a and the second copy be generated by b. Then a normal form looks like $a^{e_1} b^{f_1} a^{e_2} b^{f_2}...$ or $b^{f_1} a^{e_1} b^{f_2} a^{e_2}...$ (where the e_i and f_i are nonzero), and clearly, $\mathbb{Z} * \mathbb{Z} = \langle a, b \rangle \cong F_2$.

Now suppose we have groups G and H with isomorphic subgroups $A \leqslant G$ and $B \leqslant H$ (where $\rho : A \to B$ is a isomorphism).

Definition 5.2. The *amalgamated free product*, denoted

$$\langle G * H \mid a = \rho(a), a \in A \rangle,$$

is the quotient group of the free product $G * H$ by the normal closure of the set $\{\rho(a)a^{-1} \mid a \in A\}$.
We often write $G *_A H$ (if ρ is understood).

We can think of the amalgamated free product $G *_A H$ as the free product of G and H, glued together along the isomorphic copies of A in each group. To see this, consider the following example:

Example 5.3. Oriented graphs X and Y (shown below) have the same fundamental group F_2 (by Proposition 4.2). Suppose $\pi_1(X) = \langle a, b \rangle$ and $\pi_1(Y) = \langle c, d \rangle$, and the free product of these groups is taken with the amalgamation of the isomorphic copies of $\mathbb{Z}$ in the groups. In particular, $\langle b \rangle \cong \langle c \rangle$ via the isomorphism $\rho : b \mapsto c$. This can be seen by "gluing" the graphs X and Y together along the loop corresponding to b in X and c in Y to form graph Z. The fundamental group of the graph Z is the amalgamated product $F_2 *_{\mathbb{Z}} F_2 \cong F_3$:

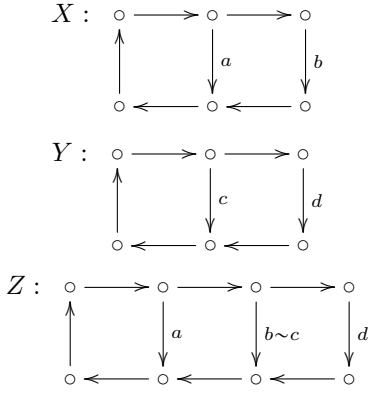

We now define an *A-normal form* for the free amalgamated product $G *_A H$. Choose a set of representatives T_A of the right cosets of A in G and a set of representatives T_B of the right cosets of $\rho(A)$ in H. Assume we have $1 \in T_A$ representing the coset A and $1 \in T_B$ representing the coset $\rho(A)$. Then we can define an A-normal form:

Definition 5.4. An *A-normal form* is a sequence $(x_0, ..., x_n)$ such that:
i) $x_0 \in A$
ii) $x_i \in T_A - \{1\}$ or $x_i \in T_B - \{1\}$ for all i, $1 \leqslant i \leqslant n$
iii) either $x_i \in T_A, x_{i+1} \in T_B$ or $x_i \in T_B, x_{i+1} \in T_A$ for all $i \geqslant 1$
(Note that we can similarly define a B-normal form.)

For any $x \in G *_A H$, we can write x as a product of factors which form an A-normal form, and then move right to left along x, assigning coset representatives for A and $\rho(A)$ and replacing appropriate elements in A with elements in $\rho(A)$ and vice versa (via ρ and ρ^{-1}). This will give us $x = x_0...x_n$, an A-normal form. It is easy to show that such a representation is unique.

Now suppose that instead of two groups, we have a single group G with two isomorphic subgroups $A, B \in G$ (where $\rho : A \to B$ is an isomorphism). Let $t \notin G$ be a new element and $\langle t \rangle$ be the cyclic group of infinite order.

Definition 5.5. The *HNN extension* of G relative to A, B, and ρ, denoted

$$\langle G, t \mid t^{-1}at = \rho(a), a \in A \rangle,$$

is the quotient group of the free product $G * \langle t \rangle$ by the normal closure of the set $\{t^{-1}at(\rho(a))^{-1} \mid a \in A\}$.
We often write G^* (if ρ is understood).

In other words, we take two isomorphic subgroups of a group G and force them to be conjugate.

We will now prove a lemma about HNN extensions that yields some basic properties about HNN extensions and that we will use in our proof of the Adian-Rabin theorem. In order to do this, we must define a *normal form* in an HNN extension. Given an HNN extension G^* of G relative to A, B, and ρ, choose a set of representatives T_A (and T_B) of the right cosets of A (and B) in G, where $1 \in T_A$ (and $1 \in T_B$) represents the coset A (and B).

Definition 5.6. A *normal form* is a sequence $g_0 t^{\varepsilon_1} g_1 ... t^{\varepsilon_n} g_n$, where each g_i represents an element of G and each ε_i is contained in $\{-1, 1\}$, such that:
i) if $\varepsilon_i = -1$, then $g_i \in T_A$
ii) if $\varepsilon_i = 1$, then $g_i \in T_B$
iii) there is no subsequence $t^{\varepsilon} 1 t^{-\varepsilon}$

Any element $x \in G^*$ has a unique representation $x = g_0 t^{\varepsilon_1} g_1 ... t^{\varepsilon_n} g_n$, where $g_0 t^{\varepsilon_1} g_1 ... t^{\varepsilon_n} g_n$ is a normal form. A formal proof of this fact is left to the reader, but can be seen by applying the following rule: write $x = g_0 t^{\varepsilon_1} g_1 ... t^{\varepsilon_n} g_n$, where each $g_i \in G$ and each $\varepsilon_i \in \{-1, 1\}$, via the canonical homomorphism $G * \langle t \rangle \to G^*$, and move from right to left along the representation, creating the coset representatives and replacing any $t^{-1}a$, $a \in A$, with $\rho(a)t^{-1}$ and any tb, $b \in B$, with $\rho^{-1}(b)t$.

This unique normal form results in the following lemma, which yields several important properties of HNN extensions as corollaries:

Lemma 5.7 (Britton's Lemma)**.** *If a word w can be expressed $w = g_0 t^{\varepsilon_1} g_1 ... t^{\varepsilon_n} g_n$, $n \geqslant 1$, with no subwords of the form $t^{-1} g_i t$, $g_i \in A$, or $t g_j t^{-1}$, $g_j \in B$, then $w \neq 1$ in G^*.*

Proof. Given $w = g_0 t^{\varepsilon_1} g_1 ... t^{\varepsilon_n} g_n$, $n \geqslant 1$, rewrite w in its normal form using the rule described above. This rule preserves the length $|w| = 2n + 1 \geqslant 3$ (since w did not violate condition (iii) of Definition 5.6). But we know this normal form is unique, and 1 has a normal form of length less than 3, so $w \neq 1$. □

Some basic properties of HNN extensions follow from Britton's lemma, including:

Corollary 5.8. *The canonical homomorphism $\phi : G \to G^*$ is injective, and thus $G \leqslant G^*$.*

Proof. Consider any nontrivial $g \in G$. $\phi(g)$ has no subwords $t^{-1}at$, $a \in A$, or tbt^{-1}, $b \in B$, so by Britton's Lemma, $\phi(g) \neq 1$. Thus, $\ker \phi = \{1\}$ and $G \hookrightarrow G^*$. □

Corollary 5.9. *If $x \in G^*$ has finite order, then x is conjugate to some element $g \in G$.*

Proof. Assume $x \neq 1$. Let $x = g_0 t^{\varepsilon_1} g_1 ... t^{\varepsilon_n} g_n$ be the unique normal formal representation of x. We have $x^m = 1$ for some $m \geqslant 2$. Thus

$$\underbrace{(g_0 t^{\varepsilon_1} g_1 ... t^{\varepsilon_n} g_n)...(g_0 t^{\varepsilon_1} g_1 ... t^{\varepsilon_n} g_n)}_{m \text{ times}} = 1$$

By Britton's Lemma, there must be some subsequence $t^{-1}at$, $a \in A$, or tbt^{-1}, $b \in B$ in this expression of x^m. However, by Britton's lemma, we know that this subsequence cannot occur in the expression $g_0 t^{\varepsilon_1} g_1 ... t^{\varepsilon_n} g_n$. Thus, we must have $t^{\varepsilon_n} g_n g_0 t^{\varepsilon_1}$ satisfying this condition.

Without loss of generality, suppose we have $t^{\varepsilon_n} g_n g_0 t^{\varepsilon_0} = t^{-1}at$. Thus, we have $\varepsilon_n = -1$, $g_n g_0 \in A$, and $\varepsilon_1 = 1$. Let $g_n g_0 = a \in A$, where $\rho(a) = t^{-1}at = b$, $b \in B$. We can now conjugate x by g_0 and then by t to obtain:

$$x = g_0 t^{\varepsilon_1} g_1 ... g_{n-1} t^{\varepsilon_n} g_n = g_0 t g_1 ... g_{n-1} t^{-1} g_n$$
$$g_0^{-1} x g_0 = g_0^{-1} g_0 t g_1 ... g_{n-1} t^{-1} g_n g_0 = t g_1 t^{\varepsilon_2} ... g_{n-1} t^{-1} a$$
$$t^{-1} g_0^{-1} x g_0 t = t^{-1} t g_1 t^{\varepsilon_2} ... g_{n-1} t^{-1} a t = g_1 t^{\varepsilon_2} ... t^{\varepsilon_{n-1}} g_{n-1} b$$

Note that conjugation is a group automorphism, so it preserves orders of elements. Therefore, $(g_0 t)^{-1} x (g_0 t)$ also has order m. So we can repeat the application of Britton's Lemma to this word. But the word $(g_0 t)^{-1} x (g_0 t) = g_1 t^{\varepsilon_2} ... t^{\varepsilon_{n-1}} g_{n-1} b$ has shorter length than the word $x = g_0 t^{\varepsilon_1} g_1 ... t^{\varepsilon_n} g_n$. Therefore, we proceed by induction on the length of x.

Eventually, we get to a word of length 1: $y^{-1} x y = g$, where $y \in G^*$ and $g \in G$. Thus, x is conjugate to an element of G. □

Now that we have defined amalgamated free products and HNN extensions and established some basic properties, we can now generalize these notions via the fundamental group of a *graph of groups*.

6. Fundamental Groups of Graphs of Groups

Definition 6.1. A *graph of groups* $(\mathbb{G}, Y)$ is a connected graph Y, along with:
i) a vertex group G_v for each $v \in Y^0$
ii) an edge group G_e for each $e \in Y^1$, where all $G_e = G_{\bar{e}}$
iii) a set of injections $\{\alpha_e : G_e \to G_{\alpha_e} \mid e \in Y^1\}$

We now want to construct a notion of *fundamental groups of graphs of groups* that generalizes the notion of fundamental groups of graphs developed in Section 4, and similarly compose a generalized version of the Fundamental Theorem of Bass-Serre Theory (Theorem 4.4). This theorem states that for a group G acting freely on a tree X, we can identify G with the fundamental group of $G \setminus X$. We want to similarly identify any group G acting on a tree X with the fundamental group of a graph of groups $\pi_1(\mathbb{G}, G \setminus X, T)$.

We were, however, restricted by the requirement that G be free. But now that we can play with vertex and edge groups, we can loosen this requirement. We define the fundamental group of a graph of groups $(\mathbb{G}, Y)$ in such a way that we are able to "quotient out" by those elements which would fix a vertex in the tree X. This inspires the following definition:

Definition 6.2. The *fundamental group* $\pi_1(\mathbb{G}, Y, T)$ of a graph of groups $(\mathbb{G}, Y)$ with respect to a maximal subtree T of Y is the quotient group of $F(\mathbb{G}, Y)$ by the normal closure of the set $\{t_e \mid e \in T^1\}$, where $F(\mathbb{G}, Y)$ is the quotient group of the

free product of all groups G_v, $v \in Y^0$, and the free group with basis $\{t_e \mid e \in Y^1\}$ by the normal closure of the sets $\{t_e^{-1}\alpha_e(g)t_e \cdot (\alpha_{\bar{e}}(g))^{-1} \mid e \in Y^1, g \in G\}$ and $\{t_e t_{\bar{e}} \mid e \in Y^1\}$.
(Note that we can think of $F(\mathbb{G}, Y)$ as a series of successive HNN extensions.)

Now we can account for the "non-freeness" of $G \curvearrowright X$. This action is not free if there exist $v \in X^0$ and $g \in G$ such that $gv = v$, i.e. if there are nontrivial stabilizer groups. We deal with the "non-freeness" by putting the stabilizers of the vertices $G_v = \{g \in G \,|\, gv = v\}$ into vertex groups and the stabilizers of the edges $G_e = \{g \in G \,|\, ge = e\}$ into edge groups.

This suggests the following proposition, a generalization of Theorem 4.4 to include actions that are not free:

Theorem 6.3 (Structure Theorem)**.** *If $G = \pi_1(\mathbb{G}, Y, T)$, then G acts without inversion of edges on a tree X such that $Y \cong G \setminus X$ and the stabilizers of the vertices and edges of the tree X are conjugate to the images (under the canonical morphism) of G_v, $v \in Y^0$, and $\alpha_e(G_e)$, $e \in Y^1$, respectively.*

Moreover, for the corresponding map $p : X \to Y$, we can lift the pair (Y, T) to $(\tilde{Y}, \tilde{T})$ such that:

i) the stabilizer of any $\tilde{v} \in \tilde{T}^0$ in G is equal to the group G_v (and the stabilizer of any $\tilde{e} \in \tilde{Y}^1$ with $\alpha(\tilde{e}) \in \tilde{T}^0$ in G is equal to the group $\alpha_e(G_e)$)

ii) if for some $\tilde{e} \in \tilde{Y}^1$ we have $\omega(\tilde{e}) \notin \tilde{T}^0$, then t_e^{-1} carries this vertex into $\tilde{T}^0$.

Conversely, if G acts on a tree X, then G is isomorphic to the fundamental group $\pi_1(\mathbb{G}, G \setminus X, T)$, where T is a maximal subtree of $G \setminus X$ and the vertex and edge groups correspond to the stabilizers of the vertices and edges of X, respectively.

A formal proof of this generalized theorem is beyond the scope of this paper, but can be found in [4].

We now demonstrate some examples of these fundamental groups of graphs of groups, beginning with the simplest case:

Example 6.4. If each vertex and edge group is trivial, our action $G \curvearrowright X$ is free and $G \cong \pi_1(G \setminus X, x)$. We see that the Structure Theorem is in fact a stronger version of Theorem 4.4.

We now want to recover amalgamated free products and HNN extensions as examples of fundamental groups of graphs of groups. To do this, we need to identify the factor graphs of these objects acting on trees. We make the following two claims:

Proposition 6.5. *Let $G = G_1 *_A G_2$. Then there exists a tree X on which G acts without inversion of edges such that $G \setminus X$ is a segment, where we define a segment as a graph with two vertices and an edge (and its inverse) connecting them. Moreover, this segment can be lifted to a segment in X such that the stabilizers in G of its vertices and edges are equal to G_1, G_2, and A, respectively.*

Proposition 6.6. *Let $G^* = \langle G, t \mid t^{-1}at = \rho(a), a \in A\rangle$. Then there exists a tree X on which G^* acts without inversion of edges such that $G^* \setminus X$ is a loop, where we define a loop as a graph with one vertex and an edge (and its inverse) from that vertex to itself. Moreover, there is a segment in X such that the stabilizers of its vertices and edges in G^* are equal to G, tGt^{-1}, and A, respectively.*

Proof. We provide a proof of Proposition 6.5, and the proof of Proposition 6.6 is very similar.

Given $G = G_1 *_A G_2$, we define the graph X as follows. Let the set of vertices X^0 be the set of left cosets G/G_1 and G/G_2. Let the set of edges X^1 be the set of left cosets G/A, where for each $gA \in X^1$, $\alpha(gA) = gG_1$ and $\omega(gA) = gG_2$. The group action $G \curvearrowright X$ will be G acting on these cosets via left multiplication. Clearly, this action is without inversion of edges.

First, we prove that X is a connected graph. Note that the vertices G_1 and G_2 are connected via edge A, so if the vertices in G/G_1 are connected and the vertices in G/G_2 are connected, then X is connected. Without loss of generality, we show that the vertices in G/G_1 are connected, i.e. that there is a path between G_1 and gG_1 for any $g \in G$. Write $g = g_0...g_n$, where $(g_0, ..., g_n)$ is the unique A-normal form of g (Definition 5.4). Then for each i, we have $g_i \in G_1, g_{i+1} \in G_2$ or $g_i \in G_2, g_{i+1} \in G_1$. If $g_i \in G_1$, then $G_1 = g_iG_1$, so the vertex $g_1...g_{i-1}g_iG_1$ is exactly the vertex $g_1...g_{i-1}G_1$. Otherwise, we have $g_i \in G_2$, and there is a path from $g_1...g_{i-1}G_1$ to $g_1...g_{i-1}G_2$ (which is equivalent to $g_1...g_{i-1}g_iG_2$) to $g_1...g_{i-1}g_iG_1$. In either case, we have $g_1...g_{i-1}G_1$ and $g_1...g_{i-1}g_iG_1$ connected, so we proceed by induction on i. Hence, X is connected.

Now, we show that X is a tree. By contradiction, suppose we have a closed, reduced path $e_0...e_m$ $(m > 0)$ in X. Without loss of generality, suppose $\alpha(e_0) = \omega(e_m) = G_1$. Note that for each i, $\alpha(e_i) \in G/G_1 \implies \omega(e_i) \in G/G_2$ and $\alpha(e_i) \in G/G_1 \implies \omega(e_i) \in G/G_2$, so we must have m odd. Set $k = \frac{m-1}{2}$. We have elements $a_0, ..., a_k \in G_1 - A$ and $b_0, ..., b_k \in G_2 - A$ such that

$$\alpha(e_1) = a_0G_2 \implies \alpha(e_2) = a_0b_0G_1 \implies ... \implies \alpha(e_m) = a_0b_0...a_kG_2$$

$$\implies \omega(e_m) = a_0b_0...a_kb_kG_1$$

But we also have $\omega(e_m) = \alpha(e_0) = G_1$, which contradicts the uniqueness of A-normal forms in $G_1 *_A G_2$. Hence, we have our tree X on which $G = G_1 *_A G_2$ acts. $G \curvearrowright X$ has two orbits of vertices: those in G/G_1 and those in G/G_2. Furthermore, these are clearly connected by the orbit of all edges G/A, so $G \setminus X$ is a segment and can be lifted to the following segment in X:

$$G_1 \overset{A}{\text{———}} G_2$$

□

Now we can easily define a graph of groups $(\mathbb{G}, Y)$ so that the fundamental group $G = \pi_1(\mathbb{G}, Y, T)$ is either an amalgamated free product or an HNN extension:

Example 6.7. Suppose Y is a segment:

$$P \overset{e}{\text{———}} Q$$

Then take $T = Y$, and G is isomorphic to the free product of G_P and G_Q amalgamated over the subgroups $\alpha_e(G_e)$ and $\alpha_{\bar{e}}(G_e)$.

Example 6.8. Suppose Y is a loop:

$$P \circlearrowright e$$

Then take $T = P$, and G is isomorphic to the HNN extension of G_P with associated subgroups $\alpha_e(G_e)$ and $\alpha_{\bar{e}}(G_e)$.

Note that given any group of graphs $(\mathbb{G}, Y)$, we can construct $G = \pi_1(\mathbb{G}, Y, T)$ via successive amalgamated products and HNN extensions. We can first determine $H = \pi_1(\mathbb{G}, T, T)$ by constructing amalgamated products of the segments in T. Then, we can think of T as contracted to this group H, and take successive HNN extensions to determine G.

Fundamental groups of graphs of groups yield some interesting extensions (to complexes, etc.) in topology, specifically in covering space theory. A discussion of this can be found here: [1]. However, we instead delve into some applications of amalgamated free groups and HNN extensions to group-theoretic decision problems.

7. The Novikov-Boone Theorem

Now that we have explored HNN extensions and amalgamated free products in the context of combinatorial group theory, we demonstrate their power in proving some very different propositions, beginning with the Novikov-Boone Theorem on the undecidability of word problems. We begin with some preliminary definitions:

Definition 7.1. Given a group G, a set X, and a set of relations R, we say G has *presentation* $\langle X \,|\, R \rangle$ if $G \cong F/N$, where F is the free group generated by X and N is the normal subgroup generated by the relations in R. Such a presentation is called $finite$ if both X and R are finite. Additionally, such a presentation is called *recursive* if R is recursively enumerable (see Definition 7.5).

Example 7.2. The dihedral group D_{2n} has presentation

$$\langle r, s \,|\, r^{2n} = 1, s^2 = 1, s^{-1}rs = 1 \rangle$$

where r can be thought of as a rotation by π/n and s can be thought of as a reflection across an axis of symmetry.
Also, the cyclic group of order n has presentation $\langle x \,|\, x^n = 1 \rangle$.

We will be interested in the *word problem for groups*: Given a group $G = \langle X \,|\, R \rangle$ and words w, w' in the generators, is $w = w'$ in G? (Or equivalently, given a word w in the generators, is $w = 1$?) Note that by word, we just mean a sequence $x_1 ... x_n$, $x_i \in X \cup X^{-1}$.

We can now state the Novikov-Boone Theorem:

Theorem 7.3 (The Novikov-Boone Theorem)**.** *There exists some finitely presented group G such that the word problem in G is undecidable, i.e. there exists a finitely presented group G for which there is no algorithm to determine whether a word in the generators represents the identity in G.*

Our proof of this theorem will rely on the following lemma:

Lemma 7.4 (Higman Embedding Theorem)**.** *A finitely generated group G can be embedded in a finitely presented group if and only if G is recursively presented.*

A proof of this theorem is omitted, but can be found in [5].

In order to prove the Novikov-Boone Theorem, we will take advantage of some concepts in computability theory. An in-depth treatment of this topic is beyond the scope of this paper, but we do provide the following definitions before beginning our proof:

Definition 7.5. A set $S \subseteq \mathbb{N}$ is *recursively enumerable* if there is an algorithm that enumerates the elements of S.

Definition 7.6. A set $S \subseteq \mathbb{N}$ is *recursive* if if there is an algorithm that halts on *all* inputs and decides whether or not the input belongs to S.

Note that recursive sets are always recursively enumerable, but recursively enumerable sets are not always recursive. Proof of the existence of a set that is recursively enumerable but not recursive is left to the reader.

We can now begin our proof of the Novikov-Boone Theorem:

Proof. Let S be a recursively enumerable but non-recursive set. We want to construct a group G_S so that the membership problem in S (which is undecidable since S is not recursive) can be reduced to the the word problem in G_S. This will show that the word problem in G_S is undecidable.

Let $F = \langle a, b \rangle$ and $F' = \langle c, d \rangle$ be two isomorphic copies of F_2. Take the free product of F and F' with amalgamation of $\langle a^{-s}ba^s \,|\, s \in S \rangle$ and $\langle c^{-s}dc^s \,|\, s \in S \rangle$ (which are clearly isomorphic via $\rho : s \mapsto s$):

$$G_S = \langle a, b, c, d \mid a^{-s}ba^s = c^{-s}dc^s, s \in S \rangle$$

Now consider the word $w = a^{-t}ba^tc^{-t}d^{-1}c^t$ in G_S:

$$w = 1 \iff a^{-t}ba^t = (c^{-t}d^{-1}c^t)^{-1} \iff a^{-t}ba^t = c^{-t}dc^t \iff t \in S$$

Hence, the membership problem in S can be reduced to the word problem in G_S, and the word problem in G_S is thus undecidable. And since S is recursively enumerable, the Higman Embedding Theorem tells us that G_S can be embedded in a finitely presented group (in which the word problem will clearly also be undecidable). □

We have only shown that there exists *some* finitely presented group G in which the word problem is undecidable. Explicit finite presentations for groups in which the word problem is undecidable have been found, and a reasonably short one (discovered by Collins and discussed in [6]) is provided for context:

$$\langle\, a, b, c, d, e, p, q, r, t, k \,|$$

$$\begin{array}{lll}
p^{10}a = ap, & pacqr = rpcaq, & ra = ar, \\
p^{10}b = bp, & p^2adq^2r = rp^2daq^2, & rb = br, \\
p^{10}c = cp, & p^3bcq^3r = rp^3cbq^3, & rc = cr, \\
p^{10}d = dp, & p^4bdq^4r = rp^4dbq^4, & rd = dr, \\
p^{10}e = ep, & p^5ceq^5r = rp^5ecaq^5, & re = er, \\
aq^{10} = qa, & p^6deq^6r = rp^6edbq^6, & pt = tp, \\
bq^{10} = qb, & p^7cdcq^7r = rp^7cdceq^7, & qt = tq, \\
cq^{10} = qc, & p^8ca^3q^8r = rp^8a^3q^8, & \\
dq^{10} = qd, & p^9da^3q^9r = rp^9a^3q^9, & \\
q^{10} = qe, & a^{-3}ta^3k = ka^{-3}ta^3 \;\rangle &
\end{array}$$

8. The Adian-Rabin Theorem

Finally, we apply the Novikov-Boone Theorem (along with further applications of amalgamated free products and HNN extensions) to prove the Adian-Rabin Theorem on the undecidability of Markov properties of finitely presented groups. We begin with a necessary definition:

Definition 8.1. A *group property* $\mathcal{P}$ is a map ϕ from the set of all groups to the set {true, false}. We require that if $G_1 \cong G_2$, then $\phi(G_1) = \phi(G_2)$. A property $\mathcal{M}$ of finitely presented groups is *Markov* if the following additional conditions are satisfied:

i) $\phi \neq 0$, i.e. there exists some finitely presented group satisfying $\mathcal{M}$

ii) there exists a finitely presented group that cannot be embedded into any finitely presented group that satisfies $\mathcal{M}$

We now provide some examples of Markov properties to preface the significance of the Adian-Rabin Theorem:

Proposition 8.2. *The property $\mathcal{M}$ of finitely presented groups such that $\mathcal{M}(G)$ is true if and only if G is abelian is a Markov group property.*

Proof. If $G_1 \cong G_2$, then clearly G_1 is abelian $\iff$ G_2 is abelian, so $\mathcal{M}$ is a group property. Furthermore, the cyclic group of order n is abelian and finitely presented (as shown in Definition 5.1), so $\mathcal{M}$ satisfies condition (i) of Definition 6.1. Furthermore, we cannot embed (via θ) a nonabelian group H (i.e. S_3 which clearly has a finite presentation) in an abelian group G: take $a, b \in H$ such that $ab \neq ba \implies \theta(ab) \neq \theta(ba) \implies \theta(a)\theta(b) \neq \theta(b)\theta(a)$, which would contradict G being abelian. Hence, $\mathcal{M}$ also satisfies condition (ii), and is Markov. □

It is easy to show that being trivial, being finite, being free, being cyclic, being simple, being solvable, and being torsion-free are also Markov group properties. This theorem will thus end up giving us many undecidable decision problems in group theory.

We can now state and prove the Adian-Rabin Theorem:

Theorem 8.3. *There is no algorithm to decide if a finitely presented group satisfies a given Markov property.*

Proof. Let $\mathcal{M}$ be a Markov property. By definition, we have some finitely presented group E satisfying $\mathcal{M}$ and some finitely presented group F that cannot be embedded in any finitely presented group satisfying $\mathcal{M}$. By the Novikov-Boone Theorem (proved in Section 7), we also have a finitely presented group H such that the word problem in H is undecidable. Consider the free product $H * F$. Since H has an undecidable word problem, $H * F$ also does. Also, since H and F both have finite presentations, $H * F$ also has a finite presentation. Let this be given by $H * F = \langle x_1, ..., x_n \,|\, R(x_1, .., x_n)\rangle$, where R is a set of relations on the generators. This proof is completed by reduction of the word problem in $H * F$ (which we know to be undecidable) to the Markov problem:

Suppose we are given a word w in the generators $x_1, ..., x_n$ of $H * F$. We will construct a finitely presented group G_w such that G_w satisfies $\mathcal{M} \iff w = 1$, thus proving the theorem. $G_w = E$ if $w = 1$ and $F \hookrightarrow G_w$ otherwise will suffice.

First, let $\langle s_0 \rangle$ be the infinite cyclic group generated by s_0, and consider the free product $G = (H * F) * \langle s_0 \rangle$. Define $s_i = s_0 x_i$ for $1 \leqslant i \leqslant n$, and we have that these

$s_0, ..., s_n$ now generate G:

$$G = \langle s_0, ..., s_n | R(x_1, ..., x_n) \rangle = \langle s_0, ..., s_n | R(s_0^{-1}s_1, ..., s_0^{-1}s_n) \rangle$$

Take $n+1$ successive HNN extensions of G relative to associated subgroups $\langle s_0, ..., s_n \rangle$ and $\langle s_0^2, ..., s_n^2 \rangle$, with the $n+1$ corresponding stable letters $t_0, ..., t_n$. We obtain the group:

$$G' = \langle s_0, ..., s_n, t_0, ..., t_n | R(s_0^{-1}s_1, ..., s_0^{-1}s_n), t_i s_i t_i^{-1} = s_i^2, 1 \leqslant i \leqslant n \rangle$$

Since our t_i are not in any of our associated subgroups, Britton's Lemma (Lemma 5.6) applied to our successive HNN extensions tells us that the subgroups $\langle t_0, ..., t_n \rangle$ and $\langle t_0^2, ..., t_n^2 \rangle$ are free on those generators. Therefore, the subgroups $\langle t_0, ..., t_n \rangle$ and $\langle t_0^2, ..., t_n^2 \rangle$ are isomorphic under $\rho : t_i \mapsto t_i^2$. Since ρ defines an isomorphism of these subgroups of G', we can take an HNN extension of G' (with these associated subgroups and with stable letter u), given by:

$$G'' = \langle s_0, ..., s_n, t_0, ..., t_n, u \mid R(s_0^{-1}s1, ..., s_0^{-1}), t_i s_i t_i^{-1} = s_i^2, u t_i u^{-1} = t_i^2, 1 \leqslant i \leqslant n \rangle$$

Now, let G_0 be the infinite cyclic group:

$$G_0 = \langle a \rangle$$

Take the HNN extension of G_0 with respect to the associated subgroups $\langle a \rangle$ and $\langle a^2 \rangle$ with stable letter b to obtain G_1:

$$G_1 = \langle a, b \mid bab^{-1} = a^2 \rangle$$

Then take the HNN extension of G_1 with associated subgroups $\langle b \rangle$ and $\langle b^2 \rangle$ with stable letter c to obtain G_2:

$$G_2 = \langle a, b, c \mid bab^{-1} = a^2, cbc^{-1} = b^2 \rangle$$

At this point we can define a penultimate group J_w, into which we claim F will inject if $w \neq 1$ and which we claim will be trivial if $w = 1$. Given a word w, define J_w as the free product of G'' and G_2, subject to the relations $u = a$ and $c = [w, s_0]$ (where $[w, s_0]$ is the commutator $ws_0w^{-1}s_0^{-1}$ of w and s_0).

First, we claim that if $w = 1$, then this group is trivial:

$$w = 1 \implies [w, s_0] = 1 \implies c = 1 \implies b = 1 \implies a = 1 \implies u = 1$$
$$\implies t_i = 1 \forall i \implies s_i = 1 \forall i \implies J_w = \langle 1 \rangle$$

Next, we show that $F \hookrightarrow J_w$ otherwise. If $w \neq 1$ then the commutator $[w, s_0]$ has infinite order in G, and thus in G''. But c is a stable letter with respect to G_1, so $u = a$ and $[w, s_0] = c$ freely generate $\langle u, [w, s_0] \rangle$. Therefore, $\langle u, [w, s_0] \rangle$ and $\langle a, c \rangle$ are isomorphic copies of F_2 in G'' and G_2, respectively, and we can think of J_w as the the amalgamated free product of G'' and G_2, with associated subgroups $\langle u, [w, s_0] \rangle$ and $\langle a, c \rangle$. Hence, we have the following sequence of injections:

$$F \hookrightarrow H * F \hookrightarrow G \hookrightarrow G' \hookrightarrow G'' \hookrightarrow J_w$$

But by assumption, F cannot be embedded in any finitely presented group satisfying $\mathcal{M}$, so J_w must not satisfy $\mathcal{M}$.

Now that we have a finitely presented J_w that is trivial if $w = 1$ and contains a group F that does not satisfy $\mathcal{M}$ if $w \neq 1$, we simply define G_w as the free product of J_w and E (which we know *does* satisfy $\mathcal{M}$):

$$G_w = J_w * E$$

If $w = 1$, $G_w = E \implies G_w$ satisfies $\mathcal{M}$. If $w \neq 1$, $G_w = J_w * E \implies F \hookrightarrow J_w \hookrightarrow G \implies G_w$ does not satisfy $\mathcal{M}$. At this point, we have reduced the word problem for finitely presented groups to the Markov property problem for finitely presented groups. Thus, there is no algorithm to decide if a finitely presented group satisfies a given Markov property. □

The Adian-Rabin Theorem thus tells us that, given a finite group presentation, there is no algorithm to determine whether or not that group is abelian, trivial, finite, free, cyclic, simple, solvable, or torsion-free.

The Adian-Rabin Theorem also yields the undecidability of Dehn's proposed group isomorphism problem (for finitely presented groups) as a corollary:

Corollary 8.4. *There is no algorithm to determine whether or not two finite group presentations represent the same group.*

Proof. If there were such an algorithm, we could compare any finitely presented group G to the trivial group, but this contradicts the Adian-Rabin theorem. Therefore, the group isomorphism problem for finitely presented groups is undecidable. □

We will conclude with a brief application of these group theoretic decision problems to a different area of mathematics, topology, to demonstrate the scope of such studies. In particular, the Adian-Rabin theorem has an important consequence in manifold theory. There is the following theorem in algebraic topology:

Theorem 8.5. *For every $n \geqslant 4$, every finitely presented group G is isomorphic to the fundamental group of a closed n-manifold M. Furthermore, given a finite presentation of G, there is an algorithm to compute a representation of such an M, say, as a simplicial complex.*

Therefore, if there were algorithm to determine whether or not a given n-manifold is simply connected, then given a finitely presented group G, we could compute an n-manifold with fundamental group G, decide if that manifold is simply connected, and consequently decide if G is trivial. This would contradict the Adian-Rabin Theorem, and so we get the following corollary of Theorem 8.5:

Corollary 8.6. *Let $n \geqslant 4$. There is no algorithm to decide, given an n-manifold M as a simplicial complex, whether M is simply connected.*

Acknowledgments. It is a pleasure to thank my mentor, Karl Schaefer, for guiding me to various sources, providing invaluable insights and examples, and patiently revising my paper. I would also like to thank all of the participating professors for their time, especially Peter May for putting this wonderful experience together.

References

[1] Oleg Bogopolski. Introduction to Combinatorial Group Theory.
[2] File:Cayley Graph of F2.svg. Wikimedia Commons. `https://commons.wikimedia.org/wiki/File:Cayley_graph_of_F2.svg`
[3] History topic: Word problems for groups. `http://www-history.mcs.st-andrews.ac.uk/PrintHT/Word_problems.html`
[4] Jean-Pierre Serre. Trees.

[5] Stål Aanderaa. A proof of Higman's embedding theorem, using Britton extensions of groups. https://www.duo.uio.no/handle/10852/44194
[6] Donald J. Collins. A simple presentation of a group with unsolvable word problem. https://projecteuclid.org/download/pdf_1/euclid.ijm/1256044631
[7] Amalgamated Products and HNN Extensions. Mathematical Notes. https://chiasme.wordpress.com/2014/05/28/amalgamated-products-and-hnn-extensions/