On Groups Whose Word Problem is Solved by a Counter Automaton

Gillian Z. Elston

Department of Mathematics, Hofstra University, Hempstead, NY 11549

Gretchen Ostheimer

Department of Computer Science, Hofstra University, Hempstead, NY 11549

Abstract

We prove that a group G has a word problem that is accepted by a deterministic counter automaton with a weak inverse property if and only if G is virtually abelian. We extend this result to larger classes of groups by considering a generalization of finite state automata, counter automata and pushdown automata. Natural corollaries of our general result include a restricted version of Herbst's classification of groups for which the word problem is a one counter language and a new classification of automata that accept context-free word problems.

1 Introduction

Let H be a group given by a finite presentation $\langle X|R\rangle$, and let $W(H)$ be the word problem for H; that is, let $W(H)$ be the set of words in $(X^{\pm})^*$ that represent the identity element in H. Several authors have explored the relationship between the formal language classification of $W(H)$ and the group theoretic classification of H. It is well-known, for example, that $W(H)$ is a regular language if and only if H is finite [1]. In 1985 Muller and Schupp proved that $W(H)$ is a context-free language if and only if H has a free subgroup of finite index [2,3]. In 1991 Herbst showed that $W(H)$ is accepted by a one counter automaton if and only if H has a cyclic subgroup of finite index [4]. For a summary of these and related results, see [5].

Our main result describes the kinds of groups for which the word problem is accepted by a G-automaton. G-automata are natural generalizations of the notions of finite state automata, counter automata and pushdown automata. G-automata are referred to as *generalized automata* in [6] and as *extended*

Preprint submitted to Elsevier Science 4 August 2003

automata in [7]; they are also defined implicitly in [5]. Loosely, if G is a group, a *G-automaton* over a finite alphabet X is an automaton in which each edge is labeled by an ordered pair, the first coordinate of which is an element of G and the second coordinate of which is an element of $X^{\pm}$ or the empty word. A word w over $X^{\pm}$ is *accepted* by A_G if there is a path from the initial vertex to a terminal vertex for which the second coordinate is w and the first coordinate is 1. The *inverse property* is a weakened version of the assumption that for each edge from σ to τ labeled x there is a corresponding edge from τ to σ labeled x^{-1}. Muller and Schupp refer to this latter property as being *reversible.* The inverse property is a natural assumption when studying automata that accept a word problem for a group; for example, if A_G is a deterministic G-automaton accepting a word problem and if A_G has only one terminal vertex, then A_G satisfies the inverse property. For precise definitions of all of these terms, see Section 2.

Theorem 7 *Let H be a finitely generated group. $W(H)$ is accepted by a deterministic G-automaton with the inverse property if and only if H has a finite index subgroup K such that K is isomorphic to a subgroup of G.*

In Section 5 we give an example showing that it is necessary to assume that A_G is deterministic. It remains an open question as to whether or not it is necessary to assume that A_G satisfies the inverse property.

In Section 4 we list many corollaries to Theorem 7. One corollary is a restricted version of Herbst's result concerning word problems which are one counter languages [4].

Corollary 11 *Let H be a finitely generated group. Then $W(H)$ is accepted by a deterministic one counter automaton with the inverse property if and only if H has a cyclic subgroup of finite index.*

The following corollary provides a version of Herbst's result for the broader class of counter automata.

Corollary 13 *Let H be a finitely generated group. Then $W(H)$ is accepted by a deterministic counter automaton with the inverse property if and only if H has a free abelian subgroup of finite index.*

Although our main result is a generalization of Herbst's result, our techniques are significantly different. Muller and Schupp first show that the Cayley graph of a group with context free word problem has more than one end. They then use Stallings Structure Theorem on finitely generated groups with more than one end. Herbst, in turn, uses the Muller and Schupp result. In contrast, our techniques are completely elementary. It appears that our assumption that the G-automata satisfy the inverse property finesses the need for any deep topology.

Combining our main result with that of Muller and Schupp leads to the following corollary concerning word problems which are context-free languages [2,3].

Corollary 14 *Let H be a finitely generated group. $W(H)$ is context-free if and only if there is a deterministic G-automaton A_G with the inverse property and G free such that A_G accepts $W(H)$.*

2 Notation and Definitions

Let X be a finite set. We use X^- to denote a set of formal inverses to the elements of X and we denote $X \cup X^-$ by $X^\pm$. The free monoid on $X^\pm$ is denoted by $(X^\pm)^*$, and the free group on X by $F(X)$. The empty word in $(X^\pm)^*$ is denoted by λ. Let θ be a homomorphism from $F(X)$ onto a group H. If w is an element of $(X^\pm)^*$, then we denote by $\overline{w}$ the image of w in H under composition of θ with the natural map from $(X^\pm)^*$ to $F(X)$. The *word problem* $W(H)$ is defined by

$$W(H) = \{w \in (X^\pm)^* \ : \ \overline{w} = 1\}.$$

Let G be a group. We define a *G-automaton* A_G over X to be a finite directed graph with a distinguished initial vertex, some distinguished terminal vertices, and with edges labeled by $G \times (X^\pm \bigcup \{\lambda\})$. If k is a positive integer, a $\mathbf{Z}^k$-automaton is called a *counter automaton*, and a $\mathbf{Z}$-automaton is a *one-counter automaton.*

If p is a path in A_G, the element of G which is the first component of the label of p is denoted $g(p)$, and the element of $(X^\pm)^*$ which is the second component of the label of p is denoted by $w(p)$. If p is the empty path, $g(p)$ is the identity element and $w(p)$ is the empty word. If p and q are paths such that the final vertex of p is equal to the starting vertex of q, we denote by pq the concatenation of the two paths. As the graphs are representing automata, we shall refer to the vertices from now on as *states.*

A G-automaton over X is said to *accept* a word $w \in (X^\pm)^*$ if there is a path p from the initial state to some terminal state such that $w(p) = w$ and $g(p) = 1$. In this case p is called an *accepting path.* If A_G is a G-automaton, we denote by $\mathcal{L}(A_G)$ the language of words accepted by A_G.

A G-automaton A_G is defined to be *accessible* if for every state σ, there is a path from the initial state to σ. A_G is *trim* if every state is visited along at least one accepting path. A_G is *complete* if for every state σ and every $a \in X^\pm$, there is an edge from σ labeled by a, i.e. $w(e) = a$ for some edge e from σ. A_G

is *deterministic* if there are no edges e such that $w(e)$ is the empty word, and if for each state σ and for each $x \in X^{\pm}$, there is at most one edge e leaving σ such that $w(e) = x$.

We say that a G-automaton has the *inverse property* if for every path p from terminal state σ_1 to terminal state σ_2, there exists a path q from σ_2 to σ_1 with $w(q) = (w(p))^{-1}$.

If A_G is deterministic and complete, we introduce further notation. If w is a word in $(X^{\pm})^*$ and σ is a state in A_G, we denote by $p(\sigma, w)$ the path starting at σ such that $w(p) = w$. We let $p(w)$ denote $p(\sigma, w)$ where σ is the initial state.

In the proof, it will be useful to refer to the finite automaton over X obtained from A_G by ignoring the first component of the edge labels; we call this the *underlying finite state automaton* A of A_G, and we denote by $\mathcal{L}(A)$ the language accepted by A.

3 Preliminaries

Note that if A_G is a G-automaton over X that accepts the word problem for a group H, then H must be X-generated (though, of course, G need not be). Note also that A_G can be made accessible without changing the language of words that it accepts by simply removing all states σ for which there is no path from the initial state to σ. The following proposition shows that we can also assume that A_G is complete and trim.

Proposition 1 *If A_G is an accessible, deterministic G-automaton over X such that $\mathcal{L}(A_G) = W(H)$ for some finitely generated group H then A_G is complete and trim.*

PROOF. Let σ be a state of A_G and $a \in X^{\pm}$. Since A_G is accessible, there is a path p from the initial state to σ. Let $w(p) = w$. Then $waa^{-1}w^{-1} \in W(H)$ and is therefore accepted by A_G. Since A_G is deterministic, the only path q from the initial state such that $w(q) = w$ is p. It follows that there is an edge e leaving σ such that $w(e) = a$ and that σ is visited along the accepting path $p(waa^{-1}w^{-1})$. $\square$

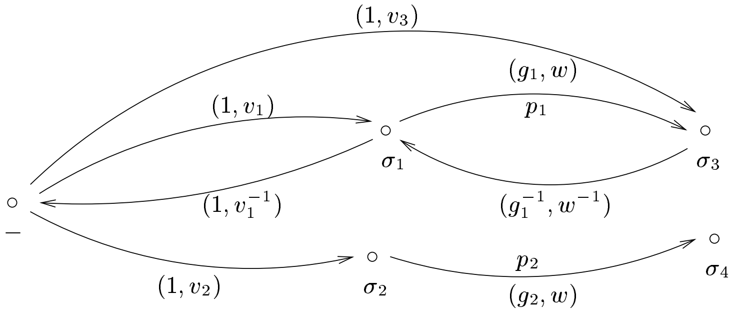

Fig. 1. Proof of Lemma 3.

4 The Word Problem and G-automata

Let G be a group. We begin by studying groups H for which the word problem is accepted by a G-automaton. Let X be a finite alphabet, and let H be a homomorphic image of $F(X)$. Let A_G be a deterministic G-automaton over X such that $\mathcal{L}(A_G) = W(H)$ and A_G satisfies the inverse property.

Note that the initial state of A_G is terminal, since the empty word is in $W(H)$. Furthermore, if σ is a terminal state of A_G, we may assume there exists a word $w \in (X^{\pm})^*$ such that $p = p(w)$ ends at σ and $g(p) = 1$: if not, then removing σ from the set of terminal states doesn't change the language accepted.

Lemma 2 *Let σ be a terminal state of A_G. Let w be a word in $(X^{\pm})^*$ such that $p(\sigma, w)$ ends in a terminal state σ'. Then $g(p(\sigma', w^{-1})) = g(p(\sigma, w))^{-1}$.*

PROOF. There exists a word u such that $p(u)$ ends at σ and $g(p(u)) = 1$. Since uww^{-1} is in $W(H)$, it follows that $g(p(\sigma', w^{-1})) = g(p(\sigma, w))^{-1}$. □

Lemma 3 *Let σ_1 and σ_2 be terminal states of A_G, $w \in (X^{\pm})^*$. If $p_1 = p(\sigma_1, w)$ ends at a terminal state, then $p_2 = p(\sigma_2, w)$ ends at a terminal state and $g(p_1) = g(p_2)$.*

PROOF. Let p_1 end at σ_3 and p_2 end at σ_4. There exist $v_1, v_2, v_3 \in \mathcal{L}(A_G) = W(H)$ such that $p(v_1)$, $p(v_2)$, and $p(v_3)$ end at σ_1, σ_2, and σ_3 respectively. By the inverse property $p(\sigma_3, w^{-1}v_1^{-1})$ ends at the initial state. Then $p_4 = p(v_3w^{-1}v_1^{-1}v_2w)$ ends at σ_4. Since $v_3w^{-1}v_1^{-1}v_2w \in W(H)$, σ_4 is a terminal state.

Now $g(p(v_1)) = g(p(v_2)) = g(p(v_3)) = 1$. Let $g_1 = g(p_1)$, and let $g_2 = g(p_2)$. By Lemma 2 $g(p(\sigma_3, w^{-1})) = g_1^{-1}$. Since $w(p_4)$ is in $W(H)$, $g(p_4) = g_1^{-1}g_2 = 1$ and $g_1 = g_2$. □

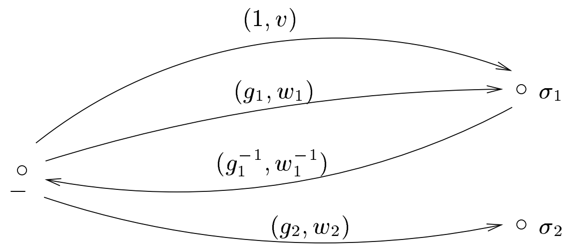

Fig. 2. Proof of Lemma 4.

Let A be the underlying finite state automaton of A_G. Let $J = \mathcal{L}(A)$, and let $K = \{\overline{w} : w \in J\}$.

Lemma 4 *Let w_1 and w_2 be words in $(X^{\pm})^*$. If $\overline{w_1} = \overline{w_2}$ for $w_1 \in J$, then $w_2 \in J$ and $g(p(w_1)) = g(p(w_2))$.*

PROOF. Let $p(w_1)$ end at σ_1 and $p(w_2)$ end at σ_2. Since $w_1 \in J$, σ_1 is a terminal state. It follows that there exists a $v \in (X^{\pm})^*$ such that $p_1 = p(v)$ ends at σ_1 and $g(p_1) = 1$, i.e. $v \in \mathcal{L}(A_G)$. Then,the inverse property implies $p(vw_1^{-1})$ ends at the initial state and $p(vw_1^{-1}w_2)$ ends at σ_2. Since $vw_1^{-1}w_2 \in W(H) = \mathcal{L}(A_G)$, σ_2 is a terminal state and $w_2 \in J$. Let $g_1 = g(p(w_1))$, and let $g_2 = g(p(w_2))$. By Lemma 2, $g(p(\sigma_1, w_1^{-1})) = g_1^{-1}$. Thus $1 = g(p(vw_1^{-1}w_2)) = g_1^{-1}g_2$ implies $g_1 = g_2$. $\square$

Lemma 5 *K is isomorphic to a subgroup of G.*

PROOF. We begin by showing that K is a subgroup of H. Let $w_1, w_2 \in J$. Since the empty word is in $W(H)$, the initial state is also a terminal state. Thus $p(w_1)$ and $p(w_2)$ both start and end at terminal states. It follows from Lemma 3 that $p(w_1w_2)$ ends at a terminal state. Thus $w_1w_2 \in J$ and J is a submonoid of $(X^{\pm})^*$, showing K is closed under multiplication.

Let $w \in J$ and let $p(w)$ end at σ. Then σ is a terminal state and $p(\sigma, w^{-1})$ ends at the initial state, also a terminal state. It follows once again from Lemma 3 that $p(w^{-1})$ ends at a terminal state and $w^{-1} \in J$. Therefore, K is closed under inverses.

Let ρ be the map from K into G that takes an element $\overline{w}$ to $g(p(w))$. ρ is well defined by Lemma 4. That ρ is a homomorphism follows directly from Lemma 3 and the fact that the initial state is also terminal. If $\overline{w} \in K$ and $\rho(\overline{w}) = 1$, then w is accepted by A_G. Thus $w \in W(H)$ and $\overline{w} = 1$. This shows ρ is injective. $\square$

Lemma 6 *K has finite index in H.*

PROOF. Let Σ be the set of states of A_G. We define a map θ from Σ to the set of right cosets of K in H as follows: $\theta(\sigma) = K\overline{w}$, where w is an element of $(X^{\pm})^*$ such that $p(w)$ ends at σ. We begin by showing that θ is well-defined. Suppose there exist paths p_1 and p_2 both starting at the initial state and ending at σ. Let q_2 be the path starting at σ such that $w(q_2) = w(p_2)^{-1}$. q_2 ends at a terminal state. Therefore p_1q_2 also ends at a terminal state, and $w(p_1q_2) = w(p_1)w(q_2) = w(p_1)w(p_2)^{-1}$ is an element of J. Therefore $K\overline{w(p_1)} = K\overline{w(p_2)}$ and θ is well-defined. Since A_G is complete, θ is onto. Since Σ is finite, the set of right cosets of K in H is also finite. $\square$

Theorem 7 *Let H be a finitely generated group. $W(H)$ is accepted by a deterministic G-automaton with the inverse property if and only if H has a finite index subgroup K such that K is isomorphic to a subgroup of G.*

PROOF. Let A_G be a deterministic G-automaton with the inverse property that accepts $W(H)$. Let A be the underlying finite state automaton of A_G, $J = \mathcal{L}(A)$, and $K = \{\overline{w} : w \in J\}$ as above. It follows from the above results that K has finite index in H and is isomorphic to a subgroup of G.

Let X be a finite alphabet, and let H be a homomorphic image of $F(X)$. Suppose that K is a finite index subgroup of H and that ρ is an embedding of K into G. We construct a G-automaton over X that accepts $W(H)$ by constructing the usual coset automaton for K in H with respect to X, with edge labels from $G \times X^{\pm}$ as described below. Let $h_1, \ldots, h_k$ be a set of right coset representatives for K in H with $h_1 = 1$. The states of A_G are the right cosets $Kh_1, \ldots, Kh_k$. Kh_1 is the initial state and the unique terminal state. There is an edge from Kh_i to Kh_m if for some x in $X^{\pm}$, $Kh_i\overline{x} = Kh_m$. That edge is labeled (g, x) where $g = \rho(h_i\overline{x}h_m^{-1})$. This defines a deterministic G-automaton over X.

To prove that $\mathcal{L}(A_G) = W(H)$ it suffices to show that if w is an element of $(X^{\pm})^*$, and if $p(w)$ ends at the state Kh_i, then $\overline{w} = \rho^{-1}(g(p(w)))h_i$. We proceed by induction on the length of w. If w is the empty word, the result is clear. Suppose that $w = w'x$, where x is an element of $X^{\pm}$. Let $g' = g(p(w'))$. Suppose that $p(w')$ ends at the state Kh_i and that $p(w)$ ends at state Kh_m. By the inductive hypothesis, we have that $\overline{w'} = \rho^{-1}(g')h_i$. By construction, there is an edge from Kh_i to Kh_m labeled (g, x), where $g = \rho(h_i\overline{x}h_m^{-1})$. We want to show that $\overline{w} = \rho^{-1}(g'g)h_m$.

$$\overline{w} = \overline{w'}\overline{x} = \rho^{-1}(g')h_i\overline{x} = \rho^{-1}(g')h_i\overline{x}h_m^{-1}h_m = \rho^{-1}(g'g)h_m.$$

Note that A_G has only one terminal state and that A_G satisfies the inverse property. $\square$

The following corollary follows immediately from Theorem 7.

Corollary 8 *Whether or not the word problem of a group H is accepted by a deterministic G-automaton satisfying the inverse property is independent of the presentation for H.*

Notice that any deterministic G-automaton with only one terminal state which accepts a word problem must satisfy the inverse property. Furthermore, examination of the proof of Theorem 7 shows that the G-automaton constructed to accept $W(H)$ has only one terminal state. For this reason, Theorem 7 could be restated as follows.

Corollary 9 *Let H be a finitely generated group. $W(H)$ is accepted by a deterministic G-automaton with only one terminal state if and only if H has a finite index subgroup K such that K is isomorphic to a subgroup of G.*

Similarly, all of the corollaries which follow could be restated by replacing the inverse property with the requirement that there be just one terminal state.

Corollary 10 *Let G be a group. Let $\mathcal{F}$ be the class of groups H for which $W(H)$ is accepted by a deterministic G-automaton A_G satisfying the inverse property. Then $\mathcal{F}$ is closed under the operations of isomorphism, finitely generated subgroups, and finite extensions.*

$\mathcal{F}$ is a family of languages as defined by Gilman in [5]. The corollary then follows from Gilman's Theorem 6.4. A more direct proof is given below.

PROOF. Closure under isomorphism is immediate. Let H be a group for which $W(H)$ is accepted by a deterministic G-automaton satisfying the inverse property. Then H has a finite index subgroup K that can be embedded in G. Let L be a finitely generated subgroup of H. Then $L \cap K$ has finite index in L, and $L \cap K$ embeds in G. Since L is finitely generated, by Theorem 7 $W(L)$ is accepted by a deterministic G-automaton satisfying the inverse property.

Let M be a finite extension of H. Then K is a finite index subgroup of M. Since H is finitely generated, so is M, and by Theorem 7 $W(M)$ is accepted by a deterministic G-automaton satisfying the inverse property. □

The following is a restricted version of Herbst's result concerning word problems which are one counter languages [4].

Corollary 11 *Let H be a finitely generated group. Then $W(H)$ is accepted by a deterministic one counter automaton with the inverse property if and only if H has a cyclic subgroup of finite index.*

PROOF. Take $G = \mathbf{Z}$. □

Let $\mathcal{S}$ be class of groups. A group H is *virtually* $\mathcal{S}$ if there exists a subgroup K of finite index in H such that $K \in \mathcal{S}$. In the case where H is finitely generated and $\mathcal{S}$ is closed under the operation of taking finitely generated subgroups, H is virtually $\mathcal{S}$ if and only if there exists a normal subgroup K of finite index in H such that $K \in \mathcal{S}$. We define an $\mathcal{S}$-automaton to be a G-automaton for some group $G \in \mathcal{S}$.

The following corollary follows immediately from Theorem 7.

Corollary 12 *Let $\mathcal{S}$ be a class of groups that is closed under the operation of taking finitely generated subgroups. Let H be a finitely generated group. $W(H)$ is accepted by a deterministic $\mathcal{S}$-automaton with the inverse property if and only if H is virtually $\mathcal{S}$.*

Corollary 13 *Let H be a finitely generated group. $W(H)$ is accepted by a deterministic counter automaton with the inverse property if and only if H has a free abelian subgroup of finite index.*

PROOF. This follows from Corollary 12 with $\mathcal{S}$ the class of free abelian groups. □

Corollary 14 *Let H be a finitely generated group. $W(H)$ is context-free if and only if there exists a free group G and a deterministic G-automaton with the inverse property such that $\mathcal{L}(A_G) = W(H)$.*

PROOF. By Corollary 12 there exists a deterministic G-automaton, with G free, satisfying the inverse property and accepting $W(H)$ if and only if H is virtually free. Muller and Schupp show that H is virtually free if and only if its word problem is context free [2,3]. □

5 A Counterexample

In this section we show that determinism is a necessary hypothesis of Theorem 7 by giving an example of groups G and H and a nondeterministic G-automaton A_G such that A_G satisfies the inverse property and accepts $W(H)$, but H does *not* have a finite-index subgroup that can be embedded in G. In this respect general G-automata accepting word problems differ from finite state automata, one-counter automata and pushdown automata that do so since in the latter settings nondeterministic automata are no more powerful than deterministic automata [2–4]. (Dassow and Mitrana showed that there exist context-free languages which are *not* word problems that cannot be accepted by deterministic G-automata [7].)

Let $X = \{x, y, z\}$, let $F = F(X)$, the free group on three generators, and let $G = F \times F$. Let $H = \langle X | x^{-1}y^{-1}xy, x^{-1}z^{-1}xz, y^{-1}z^{-1}yz, \rangle$, the free abelian group on three generators. Note that H does not have a finite index subgroup that can be embedded in G. We now construct a G-automaton A_G accepting $W(H)$.

Intuitively, we will use the edges of A_G to mimic the process of reading a word and applying the relations of H. There will be edges which correspond to

- reading a letter, and positioning the cursor to the right of the letter just read,
- moving the cursor left or right by one letter in the word that has been read so far,
- inserting a relation of H or its inverse, and positioning the cursor to the right of the inserted letters.

In this section we will distinguish between elements of $(X^{\pm})^*$ and elements of F. Let τ be the monoid homomorphism from $(X^{\pm})^*$ to F. As in the previous sections, let θ be the natural group homomorphism from F to H, and for $w \in (X^{\pm})^*$, let $\overline{w}$ represent the image of w in H, so $\overline{w} = \theta(\tau(w))$. If $u_1, u_2, \ldots, u_n$ are elements of $X^{\pm}$, and if $w = u_1 u_2 \ldots u_n$, then $reverse(w)$ is defined to be $u_n u_{n-1} \ldots u_1$. Note that for $w_1, w_2 \in (X^{\pm})^*$ such that $\tau(w_1) = \tau(w_2)$, $\tau(reverse(w_1)) = \tau(reverse(w_2))$. Therefore, for $s \in F$, it makes sense to define $reverse(s)$ to be $\tau(reverse(w))$, where w is some word in $(X^{\pm})^*$ such that $\tau(w) = s$. We will construct A_G in such a way that there is a path in A_G labeled $((s_1, s_2), w)$ if and only if $\overline{w} = \theta(s_1 reverse(s_2))$.

A_G has just one state, which is both the initial state and the terminal state. For each $u \in X^{\pm}$, there is a loop labeled $((\tau(u), 1), u)$; these edges mimic the action of reading a letter and positioning the cursor to the right of the letter just read. For each $u \in X^{\pm}$, there is also a loop labeled $((\tau(u), \tau(u^{-1})), \lambda)$, where λ represents the empty word. These edges mimic the action of moving the cursor right across a u or left across a u^{-1}. Note that if an edge labeled $((\tau(u), \tau(u^{-1})), \lambda)$ is traversed when the cursor is neither to the right of a u^{-1} nor to the left of a u, then following this edge has the effect of inserting uu^{-1} at the position of the cursor, and repositioning the cursor between u and u^{-1}. Finally, for each of the three relations $r = x^{-1}y^{-1}xy, x^{-1}z^{-1}xz, y^{-1}z^{-1}yz$ in the presentation for H, there is a loop labeled $((\tau(r), 1), \lambda)$ and another loop labeled $((\tau(r^{-1}), 1), \lambda)$; these loops mimic the action of inserting r or its inverse at the position of the cursor and repositioning the cursor to the right of the inserted letters.

Proposition 15 *There is a path in A_G labeled $((s_1, s_2), w)$ if and only if $\overline{w} = \theta(s_1 reverse(s_2))$.*

PROOF. For the first direction, we proceed by induction on the length of the path. When the path is empty, the conclusion is obvious. Suppose that there is a path p labeled $((s_1, s_2), w)$. Suppose by the inductive hypothesis that $\overline{w} = \theta(s_1 reverse(s_2))$. Suppose that e is an edge labeled $((t_1, t_2), u)$. To prove that $\overline{wu} = \theta(s_1 t_1 reverse(s_2 t_2))$, we consider the three kinds of loops separately.

First suppose that e is labeled $((\tau(u), \tau(u^{-1})), \lambda)$, where u is an element of X. Then

$$\begin{aligned}\theta(s_1\tau(u)reverse(s_2\tau(u^{-1}))) &= \theta(s_1\tau(u)\tau(u^{-1})reverse(s_2)) \\ &= \theta(s_1 reverse(s_2)) = \overline{w}.\end{aligned}$$

Next suppose that e is labeled $((\tau(r), 1), \lambda)$, where $r \in (X^{\pm})^*$ is a relation of H or its inverse. Then

$$\theta(s_1\tau(r)reverse(s_2)) = \theta(s_1 reverse(s_2)) = \overline{w}.$$

Finally suppose that e is labeled $((\tau(u), 1), u)$, where u is an element of X. Since H is abelian,

$$\theta(s_1\tau(u)reverse(s_2)) = \theta(s_1)\theta(reverse(s_2))\theta(\tau(u)) = \overline{wu}.$$

(It is possible to construct nonabelian counterexamples by introducing another state, and thereby forcing the word to be read completely before applying relations, but for our purposes, the abelian counterexample suffices.) This completes the first direction of the proof.

For the converse, suppose that s_1 and s_2 are elements of F, that w is an element of $(X^{\pm})^*$, and that $\overline{w} = \theta(s_1 reverse(s_2))$. We will construct a path labeled $((s_1, s_2), w)$. First, follow edges of the type $((\tau(u), 1), u)$ where u is in $X^{\pm}$ to form a path labeled $((\tau(w), 1), w)$. Intuitively, we have read the word w, and applied the free reductions to get a reduced word v such that $\tau(v) = \tau(w)$, and we have positioned the cursor the right of v. Let v_i be the freely reduced word representing s_i. Then v can be rewritten as $v_1 reverse(v_2)$ by inserting the relations or their inverses at suitably chosen locations in v and by applying free reductions $uu^{-1} \to \lambda$ as needed. We mimic each such rewriting step by following an appropriate sequence of edges: first follow edges of the form $((\tau(u), \tau(u^{-1})), \lambda)$ to mimic the positioning of the cursor to the desired location; then follow an edge of the form $((\tau(r), 1), \lambda)$ to mimic the insertion of a relation or its inverse at the current location. (Mimicking free reduction to the left and right of the current location happens automatically since the first coordinate of an edge label is an element of $F \times F$.) The path we have constructed is labeled $((s_1, s_2), w)$.

Thus, we have proved the proposition. $\square$

The fact that A_G accepts $W(H)$ follows immediately. Note that A_G satisfies the inverse property.

6 Open Problem

We have assumed that our automaton satisfies the inverse property, but we have not been able to establish that this hypothesis is necessary. Does there exist a finitely presented group $H = \langle X|R \rangle$ and a deterministic G-automaton A_G accepting $W(H)$ such that H does *not* have a finite index subgroup K that can be embedded in G?

7 Acknowledgments

We would like to thank Bob Gilman for suggesting this problem to us and for helpful conversations along the way.

References

[1] A. Anisimov, F. Seifert, Zur algebraischen charateristik der durch kontext-freie sprachen definierten gruppen, Elektronische Informationsverarbeitung und Kybernetik 11 (1975) 675–702.

[2] D. Muller, P. Schupp, Groups, the theory of ends and context-free languages, J. Computer and System Sciences 26 (1983) 295–310.

[3] D. Muller, P. Schupp, The theory of ends, pushdown automata, and second order logic, Theoretical Computer Science 37 (1985) 51–75.

[4] T. Herbst, On subclass of context-free groups, Theoretical Informatics and Applications 25 (1991) 255–272.

[5] R. H. Gilman, Formal languages and infinite groups, in: G. B. et. al. (Ed.), Geometric and Computational Perspectives on Infinite Groups, Vol. 25 of DIMACS Series in Discrete Mathematics and Computer Science, American Mathematical Society, Providence, RI, 1996, pp. 27–51.

[6] S. Eilenberg, Automata, Languages and Machines, Academic Press, New York, 1974.

[7] J. Dassow, V. Mitrana, Finite automata over free groups, IJAC 10 (6) (2000) 725–737.