AN AUTOMATA THEORETIC APPROACH TO THE GENERALIZED WORD PROBLEM IN GRAPHS OF GROUPS

MARKUS LOHREY AND BENJAMIN STEINBERG

ABSTRACT. We give a simpler proof using automata theory of a recent result of Kapovich, Weidmann and Myasnikov according to which so-called benign graphs of groups preserve decidability of the generalized word problem. These include graphs of groups in which edge groups are polycyclic-by-finite and vertex groups are either locally quasiconvex hyperbolic or polycyclic-by-finite and so in particular chordal graph groups (right-angled Artin groups).

1. INTRODUCTION

The *generalized word problem* is one the classical decision problems in group theory. For a finitely generated (f.g.) group G, the generalized word problem for G asks given as input elements $g, g_1, \ldots, g_n$ (represented by words over some given generating set of G) whether g belongs to the subgroup generated by $g_1, \ldots, g_n$. Examples of groups with decidable generalized word problem are f.g. free groups (see for instance [15]), polycyclic groups [1, 7], and f.g. metabelian groups [11, 12]. Moreover, every subgroup separable finitely presented group has a decidable generalized word problem. Mikhailova [8] proved that if the generalized word problem is decidable in G_1 and G_2 then the same holds for the free product $G_1 * G_2$. On the other hand, Mikhailova also proved that the direct product of two free groups of rank 2 has an undecidable generalized word problem [9]. The same was shown by Rips [10] for certain hyperbolic groups, see [17] for refinements of Rip's construction. Free solvable groups of rank 2 and derived length at least 3 also have undecidable generalized word problem [16]. It should be noted that in these undecidability results the f.g. subgroup for which membership is being asked is fixed, i.e., a fixed f.g. subgroup H of the ambient group G is constructed such that it is undecidable whether a given element of G belongs to H. On the other hand, all the decidability results mentioned above are *uniform* in the sense that the f.g. subgroup is part of the input. In order to make this distinction clear, we will often use the term "uniform generalized word problem" in the sequel.

Date: May 26, 2009.

The authors would like to acknowledge the support of DFG Mercator program. The second author is also supported by an NSERC grant.

The starting point of our work is a recent result of Kapovich, Weidmann and Myasnikov [5], which provides a condition for a graph of groups $\mathbb{G}$ that implies decidability of the uniform generalized word problem for the fundamental group $\pi_1(\mathbb{G})$. These conditions are quite technical (see the definition of a benign graph of groups in Section 4); let us just mention that (of course), for every vertex group of $\mathbb{G}$, the uniform generalized word problem has to be decidable and that every edge group has be Noetherian (i.e., does not contain an infinite ascending chain of subgroups). In [5] it is shown that graphs of groups in which edge groups are polycyclic-by-finite and vertex groups are either locally quasiconvex hyperbolic or polycyclic-by-finite — and so in particular the graphs of groups representing chordal graph groups (right-angled Artin groups) — satisfy the necessary conditions.

The proof in [5] uses an extension of the Stallings folding technique [15]. The idea is to create a "folded" graph that recognizes a normal form for each element of the subgroup. The need to be able to accept each element of the subgroup with one graph is what makes the folding moves in [5] very technical. Our proof in Section 4 is based on an automaton saturation process in the style of Benois construction [2, 3] for rational subsets of free groups, see also [4, 6]. A crucial idea is that instead of looking for membership of a given group element g in a f.g. subgroup H, we check membership of 1 in the coset Hg^{-1}. This makes the whole algorithm simpler because the normal form theorem for fundamental groups of graphs of groups is simplest for elements representing 1. The ascending chain condition for edge groups is what guarantees that our saturation process eventually terminates.

2. Cosets

Our approach to the generalized word problem is to consider cosets of finitely generated subgroups, rather than finitely generated subgroups. This has the advantage that the uniform problem reduces to checking whether the identity belongs to a coset, which is often easier.

It turns out to be useful to describe cosets in a way that does not refer to which subgroup it is a coset of and whether it is a right or left coset. Such a way was considered by Schein [13].

A *coset* of a group G is a subset A of G such that $AA^{-1}A = A$. Equivalently, if we view G as a universal algebra with a ternary operation $(x, y, z) \mapsto xy^{-1}z$, then a coset is a subalgebra of G. Traditionally, cosets are required to be non-empty, but it turns out for this paper that it is convenient to also allow the empty set to be a coset. Notice that this definition of a coset is left-right dual. One can verify that a non-empty set A is a coset in this sense if and only if $A = Hg$ for some subgroup $H \leq G$ and some element $g \in G$.

Proposition 2.1. *Let A be a non-empty coset of G. Then $H = AA^{-1}$ is a subgroup of G and if $g \in A$, then $A = Hg$. Conversely, if H is a subgroup of G and $g \in G$, then Hg is a non-empty coset of G.*

Proof. Assume first that A is a non-empty coset. Since AA^{-1} is non-empty, $AA^{-1}AA^{-1} = AA^{-1}$ and $(AA^{-1})^{-1} = AA^{-1}$, it follows that $H = AA^{-1}$ is a subgroup of G. Let $g \in A$. Clearly $Hg \subseteq AA^{-1}A = A$. Conversely, if $a \in A$ then $ag^{-1} \in AA^{-1} = H$ and so $a \in Hg$.

Clearly, if H is a subgroup of G and $g \in G$, then $(Hg)(Hg)^{-1}Hg = H^3g = Hg$ and so Hg is a non-empty coset. □

The set $K(G)$ of all cosets of G is a complete lattice with respect to the inclusion ordering since it is the set of all subalgebras of G with respect to the ternary operation considered above. The maximum element of $K(G)$ is the coset G itself. The minimal non-empty elements are the singleton subsets of G, which we identify with the elements of G notationally, i.e., we write g instead of $\{g\}$. Given any subset X of G, there is a least coset A containing X, denoted $\overline{X}$, and called the coset generated by X. It can be described as the intersection of all cosets containing X. If $X = \emptyset$, then $\overline{X} = \emptyset$ and otherwise $\overline{X}$ is non-empty. We remark that if A and B are two cosets containing g and $A = Hg$, $B = Kg$, then $A \cap B = (H \cap K)g$. Note that if A is a coset of G and H is a subgroup of G, then $A \cap H$ is a coset of H (possibly empty).

Notice that if $A \subseteq B$ are cosets and $AA^{-1} = BB^{-1}$, then $A = B$. This is clear if $A = \emptyset$. Otherwise, A, B are non-empty. Suppose $g \in A \subseteq B$. Then $B = BB^{-1}g = AA^{-1}g = A$. Thus G is Noetherian (i.e., G satisfies the ascending chain condition on subgroups, or equivalently all its subgroups are finitely generated) if and only if it satisfies the ascending chain condition on cosets.

Let us say that a coset A is *finitely generated* if there is a finite set of elements $X \subseteq G$ so that $A = \overline{X}$. That is A is finitely generated as an algebra with respect to the ternary operation $(x, y, z) \mapsto xy^{-1}z$. The next proposition shows that a non-empty coset is finitely generated if and only if it is a coset of a finitely generated subgroup; the empty coset is of course finitely generated.

Proposition 2.2. *Let A be a non-empty coset of G. Then A is finitely generated if and only if $H = AA^{-1}$ is a finitely generated subgroup of G. More specifically, let $g \in A$. If $\{g_i \mid i \in I\}$ generates H as a subgroup, then $\{g_ig \mid i \in I\} \cup \{g\}$ generates the coset A and if $\{a_i \mid i \in J\}$ generates the coset A, then $\{a_ig^{-1} \mid i \in J\}$ generates H as a subgroup.*

Proof. Assume first that $\{g_i \mid i \in I\}$ generates H as a subgroup. Then clearly $g_ig \in Hg = A$ for all $i \in I$ and $g \in Hg = A$. Now if B is any coset containing all g_ig and g, then BB^{-1} contains the g_i and hence H. Thus $A = Hg \subseteq BB^{-1}B = B$. This shows that A is generated by $\{g_ig \mid i \in I\} \cup \{g\}$. Next suppose that $\{a_i \mid i \in J\}$ generates A as a coset. Then clearly, H contains the a_ig^{-1}. If K is any subgroup containing the a_ig^{-1}, then Kg contains the a_i and hence contains A. Thus $H = AA^{-1} \subseteq Kg(Kg)^{-1} = K$. This shows that the a_ig^{-1} generate H as a subgroup. □

Notice that when going from coset generators to group generators, we can choose g to be one of the a_i.

3. Dual automata and cosets

In this section we consider an automaton model for recognizing cosets of groups. If Σ is a set we use $\widetilde{\Sigma}$ for Σ together with a set of formal inverses Σ^{-1}. Then $\widetilde{\Sigma}^*$ denotes the free monoid on $\widetilde{\Sigma}$, which we view as the free monoid with involution in the natural way. The free group on Σ will be denoted $F(\Sigma)$. When convenient, we will identify $F(\Sigma)$ with the set of reduced words in $\widetilde{\Sigma}^*$ and the canonical projection $\rho\colon \widetilde{\Sigma}^* \to F(\Sigma)$ will often be thought of as freely reducing a word. Throughout this article, Σ will be assumed finite.

3.1. **Dual automata.** By a graph Γ, we mean a graph in the sense of Serre [14]. So Γ consists of a set V of vertices, E of edges, a function $\alpha\colon E \to V$ selecting the initial vertex of an edge and a fixed-point-free involution on E written $e \mapsto e^{-1}$. This involution extends to paths in the natural way. One defines the terminal vertex function $\omega\colon E \to V$ by $\omega(e) = \alpha(\overline{e})$. A *dual automaton* $\mathscr{A}$ over Σ is a 4-tuple $(\Gamma, \iota, \tau, \delta)$ where:

- $\Gamma = (V, E)$ is a graph;
- ι, τ are distinguished vertices of Γ, called the *initial* and *terminal* vertices of Γ respectively;
- $\delta\colon E \to \widetilde{\Sigma}^*$ is an involution preserving map, i.e., $\delta(e^{-1}) = \delta(e)^{-1}$.

A dual automaton will be called *literal* if $\delta\colon E \to \widetilde{\Sigma}$.

The map δ extends to paths in the obvious way. The *language of* $\mathscr{A}$, denoted $L(\mathscr{A})$, is the subset of $F(\Sigma)$ consisting of all elements w so that there is a path p from ι to τ with $\delta(p) = w$ in $F(\Sigma)$, i.e., $\rho\delta(p) = w$. We say that p is an *accepting path* for w.

Proposition 3.1. *The language of a (finite) dual automaton over Σ is a (finitely generated) coset of $F(\Sigma)$. Conversely, every (finitely generated) coset of $F(\Sigma)$ is the language of a literal (finite) dual automaton over Σ.*

Proof. Let $\mathscr{A} = (\Gamma, \iota, \tau, \delta)$ be a dual automaton over Σ. Let $L = L(\mathscr{A})$. If L is empty, then we are done, so assume it is non-empty. It is always true that $L \subseteq LL^{-1}L$. Conversely, if $w \in LL^{-1}L$ with $w = uv^{-1}z$ such that $u, v, z \in L$ and if p, q, r are paths accepting u, v, z respectively, then $pq^{-1}r$ accepts $w = uv^{-1}z$. Thus $LL^{-1}L = L$ and so L is a coset. Notice that the map $\delta\colon E \to \widetilde{\Sigma}^*$ induces a functor, also denoted δ, from the fundamental groupoid of Γ to $F(\Sigma)$. It follows immediately from the definition that if $L \neq \emptyset$, then $LL^{-1} = \delta(\pi_1(\Gamma, \iota))$ and so alternatively we can describe L as $\delta(\pi_1(\Gamma, \iota))\delta(p)$ where p is any path from ι to τ.

Next suppose that $\mathscr{A}$ is finite and assume still that $L \neq \emptyset$. Then $\pi_1(\Gamma, \iota)$ is finitely generated and so LL^{-1} is finitely generated and hence L is finitely generated by Proposition 2.2.

Conversely, let $X \subseteq F(\Sigma)$ and let $w \in F(\Sigma)$. Define a literal dual automaton by taking a bouquet of subdivided circles at a base point ι labeled by the elements of X (with the appropriate dual edges) and attach a thorn labeled by w from ι to a new vertex τ (again with the appropriate dual edges). Then the language recognized by the resulting literal dual automaton is $\langle X \rangle w$ and the automaton is finite if X is finite. □

Notice that the proof of Proposition 3.1 is effective. If G is a group generated by Σ and $\varphi\colon \widetilde{F}(\Sigma) \to G$ is the canonical morphism, then the subset of G *recognized* by the dual automaton $\mathscr{A}$ is by definition $\varphi(L(\mathscr{A}))$. It follows that a subset of G is recognized by a finite dual automaton if and only if it is empty or a coset of a finitely generated subgroup. We obtain:

Proposition 3.2. *Let G be a group generated by a finite set Σ. Then the following are equivalent:*

(1) *The uniform generalized word problem is decidable for G.*
(2) *Uniform membership is decidable in finitely generated cosets of G.*
(3) *Uniform membership is decidable in subsets of G recognized by finite dual automata over Σ.*
(4) *There is an algorithm which given a finite dual automaton $\mathscr{A}$ over Σ as input, determines whether 1 belongs to the subset of G recognized by $\mathscr{A}$.*
(5) *There is an algorithm which given a finite literal dual automaton $\mathscr{A}$ over Σ as input, determines whether 1 belongs to the subset of G recognized by $\mathscr{A}$.*
(6) *There is an algorithm which given a finitely generated coset A of G (by a generating set) determines whether $1 \in A$.*

Proof. The equivalence of the first two items is clear. The equivalence of 2 and 3 follows from Proposition 3.1. Clearly 3 implies 4 implies 5. To see that 5 implies 6, suppose $A = \overline{X}$. If $X = \emptyset$, there is nothing to prove. Otherwise, by Proposition 2.2 we can find a generating set for $H = AA^{-1}$, and $A = Ha$ where $a \in X$. The proof of Proposition 3.1 then effectively constructs a literal dual automaton recognizing A. That 6 implies 1 follows from the observation that $g \in H$ if and only if $1 \in Hg^{-1}$ and Proposition 2.2, which allows us to effectively switch between coset generators and generators of a subgroup. □

4. The generalized word problem

In this section we use dual automata to give a technically simpler proof of a result from [5] on the decidability of the uniform generalized word problem for certain graphs of groups. We do not obtain algorithmically the induced splitting of subgroup, as is done in [5].

To make clear the main idea, we first use dual automata to give a short proof that the free product of groups with decidable generalized word problem again has decidable generalized word problem, a result due to Mikhailova

[8]. Although strictly speaking this is a special case of our main result, it seems worth proving separately to isolate the key idea.

Theorem 4.1 (Mikhailova [8]). *Let G_1, G_2 be groups with decidable uniform generalized word problem. Then the free product $G_1 * G_2$ has decidable uniform generalized word problem.*

Proof. Let Σ_1, Σ_2 be disjoint generating sets for G_1 and G_2 and put $\Sigma = \Sigma_1 \cup \Sigma_2$. Then Σ is a generating set for $G = G_1 * G_2$. By a syllable of a word $w \in \widetilde{\Sigma}^*$, we mean a maximal non-empty factor of w that can be written over a single alphabet $\widetilde{\Sigma}_i$. Let $\varphi\colon F(\Sigma) \to G$ be the projection.

Suppose that $\mathscr{A}$ is a finite literal dual automaton over Σ with initial vertex ι and terminal vertex τ. We perform the following saturation procedure. Start with $\mathscr{A}_0 = \mathscr{A}$. Assume inductively $\mathscr{A}_i$ is obtained from $\mathscr{A}_{i-1}$ by adding a new edge labeled by 1 together with its inverse edge in such a way that $\varphi(L(\mathscr{A}_i)) = \varphi(L(\mathscr{A}_{i-1}))$ for $i \geq 1$, but no vertices are added. Suppose that there is a pair p, q of distinct vertices with no edge from p to q labeled by 1 and that, for some $i = 1, 2$, there is an element of Σ_i^* representing 1 in G_i accepted by the finite dual automaton over Σ_i obtained by keeping only those edges of $\mathscr{A}_i$ labeled by elements of Σ_i or by 1, where we take the initial vertex to be p and terminal vertex to be q. Then we add an edge labeled by 1 from p to q, and the corresponding inverse edge labeled by 1 from q to p, to obtain $\mathscr{A}_{i+1}$. Otherwise the algorithm halts. Clearly $\varphi(L(\mathscr{A}_i)) = \varphi(L(\mathscr{A}_{i+1}))$. This procedure can be done effectively by Proposition 3.2 since the uniform generalized word problem is decidable in G_1 and G_2. It eventually stops since we add no new vertices. Let $\mathscr{B}$ be the final automaton obtained when the algorithm terminates.

We claim that $1 \in \varphi(L(\mathscr{A}))$ if and only if 1 labels an edge from ι to τ in $\mathscr{B}$. Since $\varphi(L(\mathscr{B})) = \varphi(L(\mathscr{A}))$, trivially if there is an edge from ι to τ in $\mathscr{B}$ labeled by 1, then $1 \in \varphi(L(\mathscr{A}))$. Conversely, suppose $1 \in \varphi(L(\mathscr{A}))$ and let w be a word accepted by $\mathscr{B}$ with $\varphi(w) = 1$ having a minimum number of syllables. If $w = 1$ or has one syllable, then by construction of $\mathscr{B}$ there is an edge labeled by 1 from ι to τ. Otherwise, the normal form theorem for free products implies w has a syllable representing 1 in one of the free factors. But then by construction of $\mathscr{B}$, the part of the accepting path for w traversed by this syllable can be replaced by a single edge labeled by 1 and so there is a word with fewer syllables accepted by $\mathscr{B}$ and mapping to 1 in G. This contradiction completes the proof. $\square$

We briefly recall the definition of a graph of groups and its fundamental group; a detailed introduction can be found in [14]. A *graph of groups* $\mathbb{G} = (G, Y)$ consists of a graph Y and

(i) for each vertex $v \in V(Y)$, a group G_v;
(ii) for each edge $y \in E(Y)$, a group G_y such that $G_y = G_{y^{-1}}$;
(ii) for each edge $y \in E(Y)$, monomorphisms $\alpha_y\colon G_y \to G_{\alpha(y)}$ and $\omega_y\colon G_y \to G_{\omega(y)}$ such that $\alpha_y = \omega_{y^{-1}}$ for all $y \in E(Y)$.

We assume that the groups G_v intersect only in the identity, and that they are disjoint from the edge set $E(Y)$. For each $v \in V(Y)$, let $\langle \Sigma_v \mid R_v \rangle$ be a presentation for G_v, with the different generating sets Σ_v disjoint. Let Δ be a set containing exactly one edge from each orbit of the involution $y \mapsto y^{-1}$ on $E(Y)$; we identify $E(Y)$ and $\widetilde{\Delta}$ when convenient. Let Σ be the (disjoint) union of all the sets Σ_v and Δ. We define a group $F(G,Y)$ by the presentation

$$F(G,Y) = \langle \Sigma \mid R_v \ \ (v \in V(Y)) \, , y\omega_y(g)y^{-1} = \alpha_y(g) \ \ (y \in E(Y), g \in G_y) \ \rangle.$$

Fix a vertex $v_0 \in V(Y)$. A word in $w \in \widetilde{\Sigma}^*$ is of *cycle type at* v_0 if it is of the form $w = w_0 y_1 w_1 y_2 w_2 \dots y_n w_n$ where:

(i) $y_i \in E(Y)$ for all $1 \leq i \leq n$;
(ii) $y_1 \cdots y_n$ is a path in Y starting and ending at v_0;
(iii) $w_0 \in \widetilde{\Sigma}^*_{v_0}$;
(iv) for $1 \leq i \leq n$, $w_i \in \widetilde{\Sigma}^*_{\omega(y_i)}$.

The images in $F(G,Y)$ of the words of cycle type at v_0 form a subgroup $\pi_1(G,Y,v_0)$ of $F(G,Y)$, called the *fundamental group of* (G,Y) *at* v_0. The fundamental group of a connected graph of groups is (up to isomorphism) independent of the choice of vertex v_0.

Let $\mathbb{G} = (G,Y)$ be a graph of groups. Let v_0 be a vertex of Y and fix a spanning tree T for Y. For a vertex v, let p_v be the unique geodesic path in T from v_0 to v. The fundamental group $H = \pi_1(G,Y,v_0)$ is generated by the words of cycle type at v_0 of the form $p_{\alpha(y)} y p_{\omega(y)}^{-1}$ with $y \in \Delta \setminus T$ and $p_v x p_v^{-1}$ with $x \in \Sigma_v$. In particular, if Y is finite and each of the vertex groups is finitely generated, then the fundamental group is finitely generated.

A graph of groups $\mathbb{G}$ is *benign* [5] if the following conditions hold:

(1) For each vertex $v \in V(Y)$ and each edge $y \in E(Y)$ with $\omega(y) = v$, there is an algorithm which given a finitely generated subgroup K of G_v (in terms of a finite generating set for K given by words from $\widetilde{\Sigma}^*_v$) and an element $g \in G_v$ (via a word in $\widetilde{\Sigma}^*_v$) determines whether $Kg \cap \omega_y(G_y)$ is empty and if it is non-empty returns an element of the intersection (represented by a word in the alphabet $\widetilde{\Sigma}_v$);
(2) Each edge group G_y is Noetherian, or equivalently, all its subgroups are finitely generated;
(3) The uniform generalized word problem is decidable for each edge group G_y;
(4) For each vertex $v \in V(Y)$ and each edge $y \in E(Y)$ with $\omega(y) = v$, there is an algorithm which given a finitely generated subgroup K of G_v (represented by a finite set of words over $\widetilde{\Sigma}^*_v$ generating K) computes a finite generating set for $K \cap \omega_y(G_y)$ as words over $\widetilde{\Sigma}_v$. Note that $K \cap \omega_y(G_y)$ must be finitely generated since $\omega_y(G_y)$ is Noetherian.

It is immediate from Proposition 2.2 that 1 and 4 from the definition of a benign graph of groups are jointy equivalent to the following statement.

(5) For each vertex $v \in V(Y)$ and each edge $y \in E(Y)$ with $\omega(y) = v$, there is an algorithm which given a finitely generated coset A of G_v (in terms of a finite generating set for the coset given by words from $\widetilde{\Sigma}_v^*$) produces a finite generating set for the (possibly empty) coset $A \cap \omega_y(G_y)$ of $\omega_y(G_y)$ (represented by words in the alphabet $\widetilde{\Sigma}_v$). Again $A \cap \omega_y(G_y)$ is finitely generated since G_y is Noetherian.

We remark that in a benign graph of groups, given an element of $\omega_y(g)$ as a word w in the generators $\widetilde{\Sigma}_{\omega(y)}$ one can find effectively a word v in the generators $\widetilde{\Sigma}_{\alpha(y)}$ representing $\alpha_y(g)$. To see this assume that G_y is finitely generated by Σ_y. Then the monomorphisms α_y and ω_y can be represented by mappings $\alpha'_y \colon \Sigma_y \to \widetilde{\Sigma}^*_{\alpha(y)}$, $\omega'_y \colon \Sigma_y \to \widetilde{\Sigma}^*_{\omega(y)}$. Then, given $w \in \widetilde{\Sigma}^*_{\omega(y)}$ one enumerates all words $u \in \widetilde{\Sigma}_y^*$ until a word u with $\omega'_y(u) = w$ is found. This word u represents $g \in G_y$ and we can compute $v = \alpha'_y(u)$. We shall use this fact below without comment.

Theorem 4.2 (Kapovich, Weidmann, Myasnikov)**.** *Let $\mathbb{G} = (G, Y)$ be a finite, connected, non-empty benign graph of finitely generated groups with underlying graph Y. Then the fundamental group of $\mathbb{G}$ has decidable uniform generalized word problem if and only if every vertex group does.*

Proof. Let $\varphi \colon F(\Sigma) \to F(G, Y)$ be the projection. We retain the above notation. In particular, we continue to use H to denote the fundamental group of $\mathbb{G}$. Since each vertex group embeds into the fundamental group [14], one implication is immediate.

Suppose that K is a finitely generated subgroup of H. We assume its generators are given as words over $\widetilde{\Sigma}$ of cycle type at v_0. Let $g \in H$ be given by a word of cycle type at v_0 and construct the literal dual automaton $\mathscr{A}$ over Σ from the proof of Proposition 3.1 recognizing Kg^{-1} (as a coset of $F(G, Y)$). Then by construction $g \in K$ if and only if there is a word of cycle type at v_0 representing 1 in $F(G, Y)$ and reading a path from ι to τ. We now perform a saturation procedure to $\mathscr{A}$ to obtain a new dual automaton $\mathscr{B}$ over Σ containing $\mathscr{A}$ with the same vertex set and the same initial and terminal vertices. Moreover, $\mathscr{B}$ will have the property that $\varphi(L(\mathscr{A})) = \varphi(L(\mathscr{B}))$ and that $g \in K$ if and only if there is an edge from ι to τ in $\mathscr{B}$ labeled by 1.

The saturation procedure continues as long as one of the following steps can be performed. We assume that at each stage of the construction, we have added no new vertices and that all edges are labeled by an element of $\widetilde{\Sigma}_v^*$, for some v, or by an element of $E(Y)$.

Suppose that the automaton at the current phase of the saturation procedure is $\mathscr{A}_i$. If $\Lambda \subseteq \Sigma$ and p, q are vertices of a dual automaton $\mathscr{C}$ over Σ, denote by $\mathscr{C}(\Lambda, p, q)$ the dual automaton consisting of all edges of $\mathscr{C}$ labeled by elements of $\widetilde{\Lambda}^*$, taking as initial vertex p and as terminal vertex q.

Step 1. If there are vertices $p \neq q$ so that $1 \in \varphi(\mathscr{A}_i(\Sigma_v, p, q)) \subseteq G_v$, then add an edge $p \xrightarrow{1} q$ and an inverse edge $q \xrightarrow{1} p$ if there are not already such edges. This step can be done effectively because of our assumption that G_v has a decidable uniform generalized word problem.

Step 2. Let $p \xrightarrow{y} q$ and $p' \xrightarrow{y} q'$ be edges in $\mathscr{A}_i$ with $y \in E(Y)$ (not necessarily distinct). Let L be the coset $\varphi(\mathscr{A}_i(\Sigma_{\omega(y)}, q, q'))$ and compute, using that the graph of groups is benign, a finite generating set X for the (possibly empty) coset $L \cap \omega_y(G_y)$ represented by words over $\widetilde{\Sigma}_{\omega(y)}$. For each $x \in X$, find a word $w_x \in \widetilde{\Sigma}^*_{\alpha(y)}$ representing $\alpha_y \omega_y^{-1}(x)$ and add an edge $p \xrightarrow{w_x} p'$ and an inverse edge $p' \xrightarrow{w_x^{-1}} p$ if $\varphi(w) \notin \varphi(\mathscr{A}_i(\Sigma_{\alpha(y)}, p, p'))$ (the latter can be checked effectively, since $G_{\alpha(y)}$ has a decidable uniform generalized word problem).

This procedure is continued until none of the steps can be performed further. Clearly this procedure does not change the accepted subset of $F(G, Y)$, which is the coset $Kg^{-1} \subseteq H$.

We must show that our procedure stops. Step 1 can only be performed finitely many times since we are adding no new vertices. Since the edge groups are Noetherian, they satisfy ascending chain condition on cosets. Step 2 can only be applied if the coset $\varphi(\mathscr{A}_i(\Sigma_{\alpha(y)}, p, p')) \cap \alpha(G_y)$ is made into a bigger coset $\varphi(\mathscr{A}_{i+1}(\Sigma_{\alpha(y)}, p, p')) \cap \alpha(G_y)$ by adding the elements of X (following the notation of Step 2) written in the alphabet $\Sigma_{\alpha(y)}$. Thus Step 2 can only be applied a finite number of times.

Hence the procedure eventually terminates with a 'saturated' dual automaton $\mathscr{B}$. The following lemma is crucial.

Lemma 4.3. *Suppose that $p \xrightarrow{y} q$ and $p' \xrightarrow{y} q'$ are edges in $\mathscr{B}$ and that $w \in \widetilde{\Sigma}^*_{\omega(y)}$ labels a path from q to q' in $\mathscr{B}$ and satisfies $\varphi(w) = \omega_y(g)$. Then there is a word $u \in \widetilde{\Sigma}^*_{\alpha(y)}$ labeling a path from p to p' in $\mathscr{B}$ with $\varphi(u) = \alpha_y(g)$.*

Proof. The element $\varphi(w) = \omega_y(g)$ belongs to the coset

$$C = \varphi(\mathscr{B}(\Sigma_{\omega(y)}, q, q')) \cap \omega_y(G_y).$$

By saturation of $\mathscr{B}$ under Step 2, $\varphi(\mathscr{B}(\Sigma_{\alpha(y)}, p, p')) \cap \alpha_y(G_y)$ contains a generating set of the coset $\alpha_y(\omega_y^{-1}(C))$ (represented by words over the alphabet $\widetilde{\Sigma}_{\alpha(y)}$). Since in fact $\varphi(\mathscr{B}(\Sigma_{\alpha(y)}, p, p')) \cap \alpha_y(G_y)$ is a coset, it thus contains $\alpha_y(\omega_y^{-1}(C))$ and so there is a word $u \in \widetilde{\Sigma}^*_{\alpha(y)}$ such that $\varphi(u) = \alpha_y(g)$ and u labels a path in $\mathscr{B}$ from p to p', as required. $\square$

We now claim that $g \in K$ if and only if 1 labels an edge from ι to τ in $\mathscr{B}$. If there is an edge from ι to τ labeled by 1, then trivially $g \in K$. Conversely, if $g \in K$ then there is a word w of cycle type at v_0 reading from ι to τ in $\mathscr{A}$ (and hence in $\mathscr{B}$) with $\varphi(w) = 1$. Choose $w = w_0 y_1 w_1 y_2 w_2 \ldots y_n w_n$ satisfying (i)–(iv) with n minimal so that $\varphi(w) = 1$ and w labels a path from ι to τ in $\mathscr{B}$. If $n = 0$, then saturation under Step 1 shows that 1 labels an edge from ι to τ. Suppose $n \geq 1$. We obtain a contradiction. It follows by [14, Theorem I.11] that there exists i such that $y_{i+1} = y_i^{-1}$ and w_i represents an element $\omega_{y_i}(g_i)$ for some $g_i \in G_{y_i}$. It now follows from Lemma 4.3 that we can replace the factor $y_i w_i y_{i+1}$ by a word $w_i' \in \widetilde{\Sigma}^*_{\alpha(y_i)} = \widetilde{\Sigma}^*_{\omega(y_{i-1})} = \widetilde{\Sigma}^*_{\alpha(y_{i+2})}$ to obtain a new word $w_0 y_1 \cdots y_{i-1} w_{i-1} w_i' w_{i+1} y_{i+2} \cdots w_n$ of cycle type at v_0 labeling a path in $\mathscr{B}$ from ι to τ and mapping to 1 in $F(G, Y)$, again a contradiction to minimality of n. This completes the proof. $\square$

References

[1] J. Avenhaus and D. Wißmann. Using rewriting techniques to solve the generalized word problem in polycyclic groups. In *Proceedings of the ACM-SIGSAM 1989 International Symposium on Symbolic and Algebraic Computation*, pages 322–337. ACM Press, 1989.

[2] M. Benois. Parties rationnelles du groupe libre. *C. R. Acad. Sci. Paris*, Sér. A, 269:1188–1190, 1969.

[3] M. Benois and J. Sakarovitch. On the complexity of some extended word problems defined by cancellation rules. *Information Processing Letters*, 23(6):281–287, 1986.

[4] M. Kambites, P. V. Silva, and B. Steinberg. On the rational subset problem for groups. *Journal of Algebra*, 309(2):622–639, 2007.

[5] I. Kapovich, R. Weidmann, and A. Myasnikov. Foldings, graphs of groups and the membership problem. *International Journal of Algebra and Computation*, 15(1):95–128, 2005.

[6] M. Lohrey and B. Steinberg. The submonoid and rational subset membership problems for graph groups. *Journal of Algebra*, 320(2):728–755, 2008.

[7] A. I. Malcev. On homomorphisms onto finite groups. *American Mathematical Society Translations, Series 2*, 119:67–79, 1983. Translation from Ivanov. Gos. Ped. Inst. Ucen. Zap. 18 (1958) 49–60.

[8] K. A. Mihaĭlova. The occurrence problem for free products of groups. *Doklady Akademii Nauk SSSR*, 127:746–748, 1959.

[9] K. A. Mihaĭlova. The occurrence problem for direct products of groups. *Math. USSR Sbornik*, 70:241–251, 1966. English translation.

[10] E. Rips. Subgroups of small cancellation groups. *Bulletin of the London Mathematical Society*, 14:45–47, 1982.

[11] N. S. Romanovskiĭ. Some algorithmic problems for solvable groups. *Algebra i Logika*, 13(1):26–34, 1974.

[12] N. S. Romanovskiĭ. The occurrence problem for extensions of abelian groups by nilpotent groups. *Sibirsk. Mat. Zh.*, 21:170–174, 1980.

[13] B. Schein. Semigroups of strong subsets. *Volzhsky Matematichesky*, 4:180–186, 1966.

[14] J.-P. Serre. *Trees.* Springer, 2003.

[15] J. R. Stallings. Topology of finite graphs. *Invent. Math.*, 71(3):551–565, 1983.

[16] U. U. Umirbaev. The occurrence problem for free solvable groups. *Sibirskiĭ Fond Algebry i Logiki. Algebra i Logika*, 34(2):211–232, 243, 1995.

[17] D. T. Wise. A residually finite version of Rips's construction. *Bulletin of the London Mathematical Society*, 35(1):23–29, 2003.

Universität Leipzig, Institut für Informatik, Germany and School of Mathematics and Statistics, Carleton University, ON, Canada

E-mail address: lohrey@informatik.uni-leipzig.de, bsteinbg@math.carleton.ca