The compressed conjugacy problem in relatively hyperbolic groups

Derek Holt and Sarah Rees

5th Dec, Warwick

Abstract

We prove that the compressed conjugacy problem in a group that is hyperbolic relative to a collection of free abelian subgroups is solvable in polynomial time.

1 Introduction

We proved in [12] that the compressed word problem is solvable in polynomial time in groups that are hyperbolic relative to a collection of free abelian subgroups. Here, we extend this result to the compressed conjugacy problem for the same class of groups, that is, we prove the following result:

Theorem A. *The compressed conjugacy problem for a group that is hyperbolic relative to a collection of free abelian subgroups is solvable in polynomial time.*

In fact, in the case when the two input elements g and h are conjugate, the algorithm that we describe constructs, in polynomial time, an SLP for a corresponding conjugator; that is, an element that conjugates g to h.

The conjugacy problem $\mathtt{CP}(G)$ for a group G takes as input two words u, v over a generating set Σ. The problem is solvable if there is an algorithm that can determine for any such input whether u, v are conjugate in G, that is, whether there exists an element g for which the product gug^{-1} represents the same element of G as v. The answer `yes` or `no` is returned.

The *compressed conjugacy problem* $\mathtt{CCP}(G)$ takes as input straight line programs (SLPs), $\mathcal{G}_1$ and $\mathcal{G}_2$, which define compressed version of words u and v (called the *values* of $\mathcal{G}_1$ and $\mathcal{G}_2$), and asks whether the group elements represented by those values are conjugate in G. In the case when the answer is positive, a solution to the problem would normally be expected to compute a conjugator g (as ours

does), which would be returned as an SLP for a word representing g. Complexity is measured in terms of the sizes of the input SLPs, which can be significantly less than the sizes of their values.

It is an easy observation that the computational complexities of the compressed word and conjugacy problems for G are independent of the choice of generating set.

We give background and notation on compressed decision problems and relative hyperbolicity of groups in Section 2, and also refer to [12] for more detail. The notation of [12] and many of its arguments are used throughout this article. Within this introduction we attempt to explain in general terms our approach to the proof of Theorem A.

From now on, we denote by G the group of the theorem, by Σ a (carefully chosen) generating set for G, by $\mathcal{G}_1$ and $\mathcal{G}_2$ SLPs defining the compressed words that are input to $\texttt{CCP}(G)$, and by u and v their values, which are (standard) words over Σ.

We assume throughout this article that the group G is hyperbolic relative to a collection of free abelian subgroups H_i, as in Section 2.3.1.

The basic idea of the proof of Theorem A is that, if the lengths of the derived words $\hat{u}$ and $\hat{v}$ of u, v (as defined in Section 2.3.1 below) are both less than some constant, then we use the methods developed in [2, Section 9] to solve the problem, and otherwise we adapt for relatively hyperbolic groups the methods that are employed in [7, Section 3] to solve the (uncompressed) conjugacy problem in linear time in hyperbolic groups.

It is shown in [10, Section 6.4] that it is straightforward to adapt the proof in [7, Section 3] to solve the compressed conjugacy problem in hyperbolic groups in polynomial time. Earlier algorithms for solving the conjugacy problem, such as those described in the proof in [4, Γ.2] for hyperbolic groups and in the proof in [2, Section 9] for relatively hyperbolic groups, involve looking at all cyclic conjugates of one or both of the input words, which cannot be done in polynomial time in the compressed setting. The proof in [7, Section 3] avoids doing this, and reduces the problem to deciding whether one word is a cyclic conjugate of another, which can be done in the compressed setting (see [13, Theorem 1]).

2 Background

Our notation and definitions follow [12], which itself largely follows [10].

2.1 Words

We define a (standard) word over an alphabet X to be a string $x_0 \cdots x_{n-1}$ of symbols from X. In this article X will always be a generating set for a group G, called either Σ or $\widehat{\Sigma}$, and will be assumed inverse closed. The word $w = x_0 \cdots x_{n-1}$ is defined to have *length* n, written $|w|$. Its subword $x_i x_{i+1} \cdots x_{j-1}$ will be denoted by $w[i, j)$, following the notation of [12]. We define the concatenation uv of words $u = u_0 \cdots u_{r-1}$ and $v = v_0 \cdots v_{s-1}$ to be the word $u_0 \cdots u_{r-1} v_0 \cdots v_{s-1}$.

For words v, w over the same alphabet, we write $v = w$ if v and w are equal as strings and we write $v =_G w$ if v and w represent the same element of G.

The *cyclic conjugates* of a word $w = x_0 \cdots x_{n-1}$ are the words $x_i x_{i+1} \cdots x_{n-1} x_0 \cdots x_{i-1}$ for $i = 0, \ldots, n-1$. Of course they represent (some of the) elements of the group G that are conjugate to w.

2.2 Straight-line programs

Let X be a finite alphabet and V a finite set with $V \cap X = \emptyset$. Let $\rho : V \to (V \cup X)^*$ be a map and extend the definition of ρ to $(V \cup X)^*$ by defining $\rho(a) = a$ for all $a \in X \cup \{\epsilon\}$ and $\rho(uv) = \rho(u)\rho(v)$ for all $u, v \in (V \cup X)^*$. We define the associated binary relation $\succeq$ on V by $A \succeq B$ whenever the symbol B occurs within the string $\rho^k(A)$, for some $k \geq 0$.

We define a *straight-line program* (SLP for short) over the alphabet X to be a triple $\mathcal{G} = (V, S, \rho)$, with $S \in V$ and $\rho : V \to (V \cup X)^*$ a map such that the associated binary relation $\succeq$ on V is acyclic, that is the corresponding directed graph contains no directed cycles. The set V is called the set of *variables* of $\mathcal{G}$, and S is called the *start variable*. Where necessary, we write $V_\mathcal{G}$, $S_\mathcal{G}$, $\rho_\mathcal{G}$, rather than simply V, S, ρ. For a variable $A \in V$, the word $\rho(A)$ is called the *right-hand side* of A. We define the *size* of $\mathcal{G}$ to be the total length of all right-hand sides: $|\mathcal{G}| := \sum_{A \in V} |\rho(A)|$.

An SLP $\mathcal{G}$ is naturally associated with a context-free grammar (V, X, S, P), where P is the set of all productions $A \to \rho(A)$ with $A \in V$, and we will often use the name $\mathcal{G}$ also for this grammar. It follows from the definition of an SLP that this associated grammar derives exactly one terminal word, which we call the *value* of $\mathcal{G}$ and denote by $\mathsf{val}(\mathcal{G})$. For any variable A of $\mathcal{G}$, we define the value of A, $\mathsf{val}(A)$ (or $\mathsf{val}_\mathcal{G}(A)$) to be the terminal word derived from that variable. Note that $\mathsf{val}(S) = \mathsf{val}(\mathcal{G})$.

SLPs are used to provide succinct representations of words that contain many repeated substrings. For instance, the word $(ab)^{2^n}$ is the value of the SLP $\mathcal{G} = (\{A_0, \ldots, A_n\}, \rho, A_0)$ with $\rho(A_n) = ab$ and $\rho(A_{i-1}) = A_i A_i$ for $0 < i \leq n$. This

SLP has size $2(n+1)$.

For any variable A of an SLP $\mathcal{G}$, we can define the *restriction* of $\mathcal{G}$ to A, $\mathcal{G}_A$. That SLP has start variable A, set of variables V_A consisting of all $B \in V$ that appear within $\rho^k(A)$ for some $k \geq 0$, and map ρ_A defined to be the restriction of ρ to V_A. We note that for any $B \in V_A$, $\mathsf{val}_{\mathcal{G}}(B) = \mathsf{val}_{\mathcal{G}_A}(B)$, and in particular $\mathsf{val}(\mathcal{G}_A) = \mathsf{val}_{\mathcal{G}}(A)$.

We provide a few technical results that we need on SLPs in Section 2.4.

2.3 Relatively hyperbolic groups

2.3.1 Definition of a relatively hyperbolic group

Our definition (below) of hyperbolicity of a group G relative to a finite collection $\{H_i : i \in \Omega\}$ of subgroups is due to Osin; it is proved in [16, Theorem 1.5] that (for finitely generated groups, as in our case) this is equivalent to the definition of [3], also to Farb's definition of [8] combined with the Coset Penetration Property (see below), called strong relative hyperbolicity in [8].

We shall use a number of properties of these groups that are proved in [2], which build on results of [16].

We suppose that Σ is a finite generating set for the group G, and that $\{H_i : i \in \Omega\}$ is a finite collection of subgroups, which we call the collection of *parabolic subgroups* of G. Define $\mathcal{H} := \bigcup_{i\in\Omega}(H_i \setminus \{1\})$, and $\widehat{\Sigma} := \Sigma \cup \mathcal{H}$. We let $\Gamma := \Gamma(G, \Sigma)$ and $\widehat{\Gamma} := \widehat{\Gamma}(G, \widehat{\Sigma})$ be the Cayley graphs for G over Σ and $\widehat{\Sigma}$, respectively. (So $\widehat{\Gamma}$ has the same vertices as Γ but more edges than Γ.) We call a word over Σ (or $\widehat{\Sigma}$) *geodesic* if it labels a geodesic path in Γ (or $\widehat{\Gamma}$).

Following [2, Definition 2.5] and [16, Section 1.2], we define F to be the free product of groups

$$F := (*_{i\in\Omega} H_i) * F(\Sigma)$$

and suppose that a finite subset R of F exists whose normal closure in F is the kernel of the natural map from F to G; in that case we say that G has the *finite presentation*

$$\left\langle \Sigma \cup \bigcup_{i\in\Omega} H_i \;\middle|\; R \right\rangle$$

relative to the collection of subgroups $\{H_i : i \in \Omega\}$. Now if u is a word over $\widehat{\Sigma}$ that represents the identity in G, then u is equal within F to a product of the form

$$\prod_{j=1}^{n} f_j r_j^{\eta_j} f_j^{-1},$$

with $r_j \in R, f_j \in F$ and $\eta_j = \pm 1$ for each j. The smallest possible value of n in any such expression of this type for u is called the *relative area* of u, denoted by $\mathsf{Area}_{\mathsf{rel}}(u)$.

We say that G is *hyperbolic relative to* the collection of subgroups $\{H_i\}$ if it has a finite relative presentation as above and a constant $C \geq 0$ such that

$$\mathsf{Area}_{\mathsf{rel}}(u) \leq C|u|$$

for all words u over $\widehat{\Sigma}$ that represent the identity in G.

We note that if G is relatively hyperbolic then the graph $\widehat{\Gamma}$ is δ-hyperbolic for some δ [16, Theorem 2.53]. Note also that, by [16, Proposition 2.36], the intersection $H_i \cap H_j$ for $i \neq j$ is finite.

An *H_i–component* of a path p in $\widehat{\Gamma}$ is defined to be a non-empty subpath of p that is maximal subject to being labelled by a word in H_i^*. We call a vertex of a path p *internal* if lies in the interior of a component of p, and otherwise *non-internal*. Two components s and r (not necessarily of the same path) are *connected* if both are H_i-components for some H_i, and if the start points of both lie in the same left coset gH_i of H_i.

A path p in $\widehat{\Gamma}$ is said to *backtrack* if $p = p'srs'p''$ where s, s' are H_i–components, and the word labelling r represents an element of H_i; if no such decomposition of p exists, then p is *without backtracking*. A path p is said to *vertex backtrack* if it contains a subpath of length greater than 1 labelled by a word that represents an element of some H_i; otherwise p is said to be *without vertex backtracking*. We note that if a path does not vertex backtrack then it does not backtrack and all of its components have length 1.

We denote the start and end points of a path p in $\widehat{\Gamma}$ by p_- and p_+, respectively, and say that paths p, q in $\widehat{\Gamma}$ are *k-similar* if $\max\{d_\Gamma(p_-, q_-), d_\Gamma(p_+, q_+)\} \leq k$. The following fundamental result about k-similar paths in $\widehat{\Gamma}$, proved as [16, Theorem 3.23], is also stated as [2, Theorem 2.8].

Proposition 2.1. *[16, Theorem 3.23] (Bounded Coset Penetration Property). Let G be relatively hyperbolic, as above. For any $\lambda \geq 1$, $c \geq 0$, $k \geq 0$, there exists a constant $\mathsf{e} = \mathsf{e}(\lambda, c, k)$ such that, for any two k-similar paths p and q in $\widehat{\Gamma}$ that are (λ, c)–quasigeodesics and do not backtrack, the following conditions hold.*

(1) The sets of vertices of p and q are contained in the closed e-neighbourhoods of each other in Γ.

(2) Suppose that, for some i, s is an H_i-component of p with $d_\Gamma(s_-, s_+) > \mathsf{e}$; then there exists an H_i-component of q that is connected to s.

(3) Suppose that s and t are connected H_i-components of p and q, respectively. Then s and t are e-similar.

We define the *components* of a word $w \in \Sigma^*$ to be the nonempty subwords of w of maximal length that lie in $(\Sigma \cap H_i)^*$ for some parabolic subgroup H_i; such a subword labels a component of any path traced out by w in the Cayley graph Γ. In general, since $H_i \cap H_j$ is finite for $i \neq j$, it is possible for the end of one component in a word w to overlap the beginning of the next, where the overlapping generators lie in a finite intersection. In this paper, we shall be assuming that the parabolic subgroups are free abelian, and hence that $H_i \cap H_j$ is trivial for $i \neq j$, and so distinct components are disjoint.

Let $w := \alpha_0 u_1 \alpha_1 u_2 \cdots u_n \alpha_n$, where the u_j are its components. Then, following [2, Construction 4.1], we define the *derived word* $\hat{w} := \alpha_0 h_1 \alpha_1 h_2 \cdots h_n \alpha_n \in \widehat{\Sigma}^*$, where each h_j is the element of a parabolic subgroup represented by u_j. So the components of paths in Γ and $\widehat{\Gamma}$ labelled by w and $\hat{w}$ are labelled by the subwords u_i and h_i of w and $\hat{w}$, respectively.

2.3.2 Some properties of a relatively hyperbolic group G

Suppose that the group G is hyperbolic relative to a collection of free abelian subgroups H_i, as in Section 2.3.1. We can select our finite generating set for G, and we select such a set Σ with particular properties that are already described in [12, Section 6.1]. We do not see the need to give the details of that construction here.

The properties of our chosen generating set Σ ensure that it contains generating sets Σ_i for each parabolic subgroup H_i. Recall that $\widehat{\Gamma}$ is a δ-hyperbolic space for some constant δ. So $\widehat{\Gamma}$ is also δ'-hyperbolic for any $\delta' > \delta$, and hence we may safely assume (and sometimes do) that $\delta \geq 1$.

We use a particular normal form $\mathsf{nf}(w)$ for words $w \in \Sigma^*$ that is already defined in [12, Section 6.1]. This has the properties that, for words w in normal form (i.e. $w = \mathsf{nf}(w)$), the derived word $\hat{w}$ labels a geodesic path in $\widehat{\Gamma}$, and the components w' of w (i.e. the maximal subwords with $w' \in \Sigma_i^*$ for some i) are shortlex reduced words. Furthermore, if w represents an element of H_i for some i, then $w \in \Sigma_i^*$ (and so $|\hat{w}| = 1$). By [12, Theorems 8.1, 9.1], for a compressed word with value $w \in \Sigma^*$, we can find a compressed word representing $\mathsf{nf}(w)$ in polynomial time, and so we can solve the compressed word problem for G in polynomial time. By [12, Proposition 3.7], the set of words in normal form is the language of an asynchronous automatic structure for G [6], and hence has particular geometrical properties that are of use to us.

As remarked in [12], it is not true that arbitrary subwords of normal form words w are in normal form, but subwords w' of w such that w' consists of generators in $\Sigma \setminus \mathcal{H}$ together with complete components of w (or, equivalently, such that $\widehat{w'}$ is a subword of $\hat{w}$) are in normal form. We shall call such subwords *non-splitting*, because they do not split components of w.

We say that a word over Σ is *stable under cyclic derivation* if it does not begin and end with letters from the same parabolic subgroup H_i. A cyclic conjugate w' of a word $w \in \Sigma^*$ is called a *non-splitting cyclic conjugate* if w' is stable under cyclic derivation, or, equivalently if $\widehat{w'}$ is a subword of $\widehat{w^2}$.

2.4 Technical results for SLPs

In [12, Proposition 4.1] we listed various properties of SLPs $\mathcal{G}$ and operations on them that can be carried out in polynomial time. These include:

(i) computation of an SLP $\mathcal{G}'$ in Chomsky normal form with $\mathsf{val}(\mathcal{G}) = \mathsf{val}(\mathcal{G}')$;
(ii) computation of $|\mathsf{val}(\mathcal{G})|$;
(iii) for any $i, j \in \mathbb{Z}$, computation of an SLP $\mathcal{G}[i : j)$ with value the substring $\mathsf{val}(\mathcal{G})[i : j)$ of $\mathcal{G}$;
(iv) given a second SLP $\mathcal{H}$ over the same alphabet as $\mathcal{G}$, we can decide whether $\mathsf{val}(\mathcal{G}) = \mathsf{val}(\mathcal{H})$.

In addition to these properties we shall need the result proved in [13, Theorem 1] that, for given SLPs $\mathcal{G}$ and $\mathcal{H}$ over the same alphabet, we can decide in polynomial time whether $\mathsf{val}(\mathcal{G})$ is a substring of $\mathsf{val}(\mathcal{H})$ and, if so, determine the smallest i such that $\mathsf{val}(\mathcal{G}) = \mathsf{val}(\mathcal{H})[i : j)$ for some j.

This last result implies immediately that we can also determine in polynomial time whether $\mathsf{val}(\mathcal{G})$ is a cyclic conjugate of $\mathsf{val}(\mathcal{H})$, by testing whether it is a substring of $\mathsf{val}(\mathcal{H})^2$.

Proposition 7.1 of [12] proves that an input SLP $\mathcal{G}$ over our chosen generating set Σ of the group G can be modified in polynomial time to produce an SLP $\mathcal{G}'$ with the same value w for which every component u of w has a root (defined to be a variable A_u with value u). Now let $\mathcal{G}'_{A_u}$ be the SLP that is the restriction of $\mathcal{G}'$ to the root A_u of u, which has A_u as its start variable. Then we can modify $\mathcal{G}'$ by collapsing each $\mathcal{G}'_{A_u}$ within $\mathcal{G}'$ to a single vertex, and then introducing a new alphabet letter a_u and defining $\rho(A_u) = a_u$. The result is an SLP for $\hat{w}$ over a finite alphabet Σ' with $\Sigma \subseteq \Sigma' \subseteq \widehat{\Sigma}$.

Hence we have proved

Proposition 2.2. *Let $\mathcal{G}$ be an* SLP *over our selected generating set Σ for G, and let $w := \mathsf{val}(\mathcal{G})$. Then, in time polynomial in $|\mathcal{G}|$, we can construct an* SLP *$\mathcal{G}'$ in Chomsky normal form with value $\hat{w}$ over a finite alphabet Σ' with $\Sigma \subseteq \Sigma' \subseteq \widehat{\Sigma}$.*

We will use this result frequently, implicitly, in our proof. It will enable us, for instance, to compute features of derived words in polynomial time, such as their lengths, and to compute SLPs for substrings v of u that are defined via substrings $\hat{v}$ of $\hat{w}$. We can also decide in polynomial time whether $\hat{u}$ is a cyclic

conjugate of $\hat{v}$, for words $u, v \in \widehat{\Sigma}$.

3 Proof of the theorem

Suppose that SLPs $\mathcal{G}_1$ and $\mathcal{G}_2$ are input, with values u and v.

We describe our algorithm in terms of the words u, v, and of words over Σ and $\widehat{\Sigma}$ that are related to those, culminating with the construction of a conjugator. But the constructions within the algorithm are of SLPs that define those words, and of course it is the construction of that sequence of SLPs, and ultimately of an SLP for the conjugator, that needs to be shown to be polynomially bounded. Since our proof consists of a possibly confusing mixture of theory and the description of the algorithm itself, we shall conclude each section with a brief summary of the steps of the algorithm presented in that section.

Since the conjugacy problem is certainly solvable in G, for any constant $C \geq 0$ we can construct in a preprocessing step a lookup table that stores the solution to $\mathtt{CCP}(G)$ for all input for which the lengths of u and v are at most C, together with corresponding conjugators (as standard words). We suppose that this has been done, for an appropriate constant C.

By [12, Theorems 8.1,9.1], which together show that an SLP may be converted in polynomial time to one with value its normal form, we may assume that our input words u and v are in normal form, so that $\hat{u}$ and $\hat{v}$ label geodesic paths in $\widehat{\Gamma}$.

Following [7, Section 3.1], we define words u_1, u_2 and u_c such that $u = u_1u_2$, $u_c := u_2u_1$ is a non-splitting cyclic conjugate of u, and $|\widehat{u_1}| \leq |\widehat{u_2}| \leq |\widehat{u_1}| + 1$. The words v_1, v_2, v_c are defined similarly. We then reduce the words u_c and v_c to normal form, and replace our original u and v by these reduced words. The motivation for doing this is that, when $\widehat{u_c}$ is long enough, all of its positive powers are quasigeodesic; this will be explained in Section 3.2.

The new words $\mathsf{nf}(u_c)$ and $\mathsf{nf}(v_c)$ are conjugates in G of the original u and v, and we must also store SLPs for corresponding conjugators, since we will need these to calculate the final conjugator in the case in the case that u and v turn out to be conjugate.

Algorithmic steps:

(i) Precompute a table of conjugacy and conjugators of words of length up to a certain constant. (We shall not attempt to specify that constant here, but its value could be computed from the various constants that we shall define later in the proof, together with constants that are defined in the results proved in [2, Section 9]. Theoretically it depends only on the group presentation.)

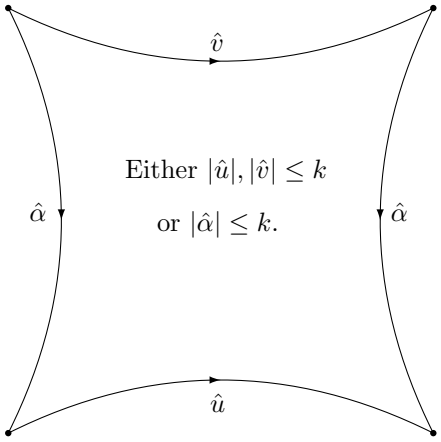

Figure 1: A minimal k-bounded conjugacy $(1,0)$ conjugacy diagram for $\hat{u}, \hat{v}$

(ii) Reduce u and v to normal form, and then compute u_c and v_c as described above. Store (as SLPs) the corresponding conjugators and reduce u_c and v_c to normal form.

3.1 Short derived words

Suppose first that the derived words $\hat{u}$ and $\hat{v}$ (after the replacement described above) both have length bounded above by the constant $17(2L'+1)/7$ with $L' := 36\delta + 2$. (Recall that δ is the hyperbolicity constant of the Cayley graph $\widehat{\Gamma}$.)

The words $\hat{u}$ and $\hat{v}$ have a bounded number of cyclic conjugates in this case and we start by calculating the normal forms of all non-splitting cyclic conjugates of u and v, and replacing u and v by cyclic conjugates in normal form such that $\hat{u}$ and $\hat{v}$ have the least possible lengths. In particular this ensures that u and v are both stable under cyclic derivation.

As in [2, Section 8], for words $\alpha, u, v \in \Sigma^*$ with $\alpha u \alpha^{-1} =_G v$, we say that the corresponding quadrilateral in $\widehat{\Gamma}$ shown in Figure 1, with paths labelled by $\hat{\alpha}$, $\hat{u}$, $\hat{\alpha}^{-1}$, $\hat{v}^{-1}$, is a *minimal conjugacy* $(1,0)$*-diagram*, if the paths labelled by $\hat{u}, \hat{v}$ and by all of their cyclic conjugates are geodesic, and if $\hat{\alpha}$ has minimal length amongst conjugators between all pairs of cyclic conjugates of $\hat{u}$ and $\hat{v}$. So if u and v are conjugate in G then there exists such a minimal conjugacy $(1,0)$-diagram in which u and v are replaced by suitable non-splitting cyclic conjugates.

We shall now consider the properties of such a minimal conjugacy diagram, assuming for now that it exists. Since $\widehat{\Gamma}$ is hyperbolic, it follows from [2, Lemma

8.2] that $(G, \widehat{\Sigma})$ has *bounded conjugacy diagrams* (BCD); that is, for some constant k, we have

$$\min\{\max\{|\hat{u}|, |\hat{v}|\}, |\hat{\alpha}|\} \leq k.$$

The fact that the paths labelled by $\hat{u}, \hat{v}$ and those labelled by their cyclic conjugates are all geodesics in $\widehat{\Gamma}$ implies immediately that the diagram is without vertex backtracking, as defined in [2, Section 9] after Corollary 9.2; that is, none of the cyclic conjugates of $\hat{u}$ and $\hat{v}$ vertex backtracks.

Now, if neither of the words u and v are parabolic (that is, they are not words over the generators of H_i for any parabolic subgroup H_i) and u and v represent conjugate elements of G, then we can apply [2, Theorem 9.13] to conclude that either the lengths of u and v are both bounded by a constant, or else there are non-splitting cyclic conjugates u' and v' of u and v such that u' and v' are conjugate by an element of G of bounded length. Since the number of possibilities for u' and v' is bounded by a constant, we can solve the problem in either of these cases, by using our precomputed list in the first case, and by exhaustive search of all possible conjugators in the second case.

Otherwise, one of the input words, say u, is parabolic, and $u \in \Sigma_i^*$ for some i. It is a basic property of relatively hyperbolic groups [2, Lemma 2.6 (iii)] that, for a parabolic subgroup H of G and $g \in G \setminus H$, the intersection $g^{-1}Hg \cap H$ is finite, and since the parabolic subgroups are free abelian in our situation, this intersection must be trivial. Hence, if v is also parabolic, then u and v are conjugate if and only if $u = v$. So we assume that the word v is not parabolic.

In what follows, we denote the path in the diagram labelled by a word w by p_w. Now, by [2, Lemma 9.14], the path p_u cannot be connected to a component of p_v. So since as we saw above $(G, \widehat{\Sigma})$ has BCD, it follows from [2, Lemma 9.16 (a)] that u has bounded length. Then by [2, Lemma 9.12] applied to p_v, the same applies to the word v, unless $\alpha u \alpha^{-1} =_G v$ with $|\hat{\alpha}| = 1$.

In that case, either $|\alpha| = 1$ and we can solve the problem by trying all possible words α, or α is parabolic. In that second case, if either of the paths p_α or $p_{\alpha^{-1}}$ is an isolated component of the conjugacy diagram (that is, is not connected to any other component) then, by [2, Lemma 9.1] applied with $I = \{p_\alpha\}$, the word α has bounded length, in which case we can again solve the problem.

So we can assume that p_α and $p_{\alpha^{-1}}$ are not isolated components. Since the components are abelian, the component p_α cannot be connected to p_u or to $p_{\alpha^{-1}}$, so it must be connected to a component of p_v. But then, by [2, Lemma 9.5] applied to the paths p_v and p_α, the paths p_α must be connected to a component of p_v that is situated at the beginning of p_v. Similarly, $p_{\alpha^{-1}}$ is connected to a component of p_v that is situated at the end of p_v. But then v is not stable under cyclic derivation, contrary to our assumption.

The discussion above justifies the following algorithmic steps for solving the conjugacy problem in the case of short derived words.

(i) Compute and reduce to normal form all of the (boundedly many) non-splitting cyclic conjugates of u. If any of these has shorter derived length than u, then replace u by that cyclic conjugate (and store the corresponding conjugator). Continue to do this until all cyclic conjugates of $\hat{u}$ have same shortest possible length. Do the same for the word v.

(ii) If u and v both have length at most a certain constant k (its value can be computed from δ and ϵ using the results in [2, Section 9]), then use the precomputed lists to test them for conjugacy in G. If they are conjugate, then return yes and a conjugator. Otherwise return no.

(iii) If at least one of u, v has length greater than k, then test all non-splitting cyclic conjugates of u and of v for conjugacy by elements α of Σ-length at most k. If any of these tests are positive then then return yes and a conjugator. Otherwise return no.

3.2 Long derived words

So we can assume now that (after the replacements of u and v by the reductions of their cyclic conjugates u_c and v_c) at least one of $|\hat{u}|$ and $|\hat{v}|$ is at least $17(2L'+1)/7$ with $L' := 36\delta + 2$. Our strategy here is to follow the proof in [7, Section 3] of the linearity of the conjugacy problem in hyperbolic groups. As remarked above, it is shown in [10, Section 6.4] that the arguments of [7] can be readily adapted to prove that the corresponding compressed conjugacy problem in hyperbolic groups is solvable in polynomial time. We are able to use them for our relatively hyperbolic group because of the hyperbolicity of the Cayley graph $\widehat{\Gamma}$ over the (infinite) generating set $\widehat{\Sigma}$. We apply the arguments in [7, Section 3] to the words $\hat{u}$, $\hat{v}$ over $\widehat{\Sigma}$. The following subsections correspond to those in [7].

3.2.1 Reduction to quasigeodesics

By [7, Lemma 3.1] (which is valid for any constant L), if $|\hat{u}| \geq 2L'+1$ then all positive powers of $\hat{u}$ define L'-local $(1, 2\delta)$-quasigeodesics in $\widehat{\Gamma}$ (that is, all of their subwords of length at most L' are $(1, 2\delta)$-quasigeodesics). Now we apply [7, Proposition 2.3] with the word w of that proposition equal to a subword of a positive power of $\hat{u}$; if that subword has length greater than L' then the last statement of the proposition asserts that its length is at most $17/7$ times the length of a geodesic between its two ends. It follows in any case that this power of $\hat{u}$ is a $(17/7, 2\delta)$-quasigeodesic.

Now if $\hat{u}$ and $\hat{v}$ are conjugate by $g \in G$, then so are all of their positive powers. It follows immediately from this (as at the end of [7, Section 3.1]) that the length of a geodesic between the two ends of $\hat{v}^n$ is at least $7|\hat{u}|n/17 - 2|g|_{\widehat{\Gamma}}$ for all $n > 0$, and hence $|\hat{v}| \geq 7|\hat{u}|/17$.

Since we are assuming that at least one of $|\hat{u}|$ and $|\hat{v}|$ is at least $17(2L'+1)/7$ it follows that, if u and v are conjugate, then $|\hat{u}|$ and $|\hat{v}|$ are both at least $2L'+1$. The argument that we applied above to $\hat{u}$ now applies to $\hat{v}$, and we deduce that all positive powers of $\hat{v}$ are also $(17/7, 2\delta)$-quasigeodesics.

To proceed further, we need the positive powers of $\hat{u}$ not to backtrack, which we can achieve as follows. If $\hat{u}^n$ backtracks for some $n > 1$, then some subword $\hat{\xi}$ of it of length greater than 1 must reduce to a parabolic element of length 1; choose $\hat{\xi}$ of maximal length with that property. Since $\hat{u}^n$ is a $(17/7, 2\delta)$-quasigeodesic, the subword $\hat{\xi}$ has length at most $17/7 + 2\delta \leq 36\delta + 2 = L'$ (recall that we are assuming that $\delta \geq 1$), and so (as an L'-local $(1, 2\delta)$-quasigeodesic) has length at most $1 + 2\delta$. Then, because $\hat{u}$ is geodesic, $\hat{\xi}$ must consist of a non-trivial suffix of $\hat{u}$ followed by a non-trivial prefix of $\hat{u}$. Note that two such subwords ξ of $\hat{u}^n$ must be separated within $\hat{u}^n$ by a subword of length at least $L' - 2\delta - 1$. (So $\hat{\xi}$ is actually a subword of $\hat{u}^2$, and appears $n-1$ times within $\hat{u}^n$.) Define $L := L' - 2\delta = 34\delta + 2$; note that this same constant L appears in [7]. After replacing each occurrence of $\hat{\xi}$ in $\hat{u}^n$ by an element of $\mathcal{H}$, the resulting word is a L-local $(1, 2\delta)$-quasigeodesic that does not backtrack.

Now u has a non-splitting cyclic conjugate that has ξ as a subword. Let u' be the word formed from that cyclic conjugate by replacing the subword ξ by $\mathsf{nf}(\xi)$. We redefine u to be this word u' (and store an SLP for the associated conjugator, which we will need later). Note that now u is stable under cyclic derivation, and so $\widehat{u^n} = \hat{u}^n$ for all $n > 0$. Now, for all $n > 0$, $\widehat{u^n}$ is a $(17/7, 2\delta)$-quasigeodesic that does not backtrack. We carry out the same process if necessary for the word v. Note that this step has reduced the lengths of $\hat{u}$ and $\hat{v}$ by at most 2δ, so they both now have length at least $2L + 1$.

By the Bounded Coset Penetration Property [16, Theorem 3.23], there is a constant ϵ such that, for any word $w \in \Sigma^*$ for which $\hat{w}$ is geodesic over $\widehat{\Sigma}$ and $w =_G u^n$, all non-internal vertices of the path in Γ starting at the origin and labelled by u^n are at Γ-distance at most ϵ from a non-internal vertex in the corresponding path labelled by $\hat{w}$, and vice versa.

Algorithmic steps:

(i) If either $|\hat{u}|$ or $|\hat{v}|$ is less than $2L'+1$ then return no.

(ii) Find all of the boundedly many non-splitting suffixes β and prefixes α of u with $|\hat{\beta}| + |\hat{\alpha}| \leq 1 + 2\delta$. For each of these nf-reduce $\beta\alpha$ and thereby find the longest non-splitting subword ξ of u^2 such that $\mathsf{nf}(\xi)$ is a parabolic word.

(iii) Compute a suitable replacement u' of u as described above together with the corresponding conjugator as SLP (stored for later).

3.2.2 Reduction to an nf-straight power

We say that a word w is nf-*straight* if all positive powers w^n of w are nf-reduced. (These corresponds to the *short-lex straight* words in [7, Section 3.2].) Note that such words are necessarily stable under cyclic derivation. We know from [12] that the set of nf-reduced words is regular, and we can use the associated finite state automaton to test compressed words for being nf-straight in polynomial time.

We say that a finite or infinite path in Γ is nf-straight if the labels of all of its finite subpaths labelled by non-splitting subwords are nf-straight.

Our argument now is basically that of [7, Section 3.2], which is itself based on an argument of Delzant. But we need some adjustment and so we give our own account. For any path γ in Γ, we define $\Delta(\gamma)$ to be the set of paths γ' in Γ for which each non-internal vertex of γ' is within Γ-distance ϵ of some non-internal vertex of γ.

Let p_n be the vertex of Γ labelled by u^n for $n \in \mathbb{Z}$ (so p_0 is the base point of Γ), and let γ_1 be the two-way-infinite path in Γ containing each p_n such that the subpath from p_n to p_{n+1} is labelled by u for each n. (The corresponding path is called w in [7].)

We let Π_0 be the set of two-way-infinite paths γ_2 in Γ with nf-straight labels for which $\gamma_2 \in \Delta(\gamma_1)$ and $\gamma_1 \in \Delta(\gamma_2)$. Now let Π be the subset of Π_0 consisting of paths going in the same direction as γ_1. We shall prove first that $|\Pi|$ is bounded above, and second that Π is non-empty.

For a path $\gamma_2 \in \Pi$ and each $n \in \mathbb{Z}$, choose a specific vertex r_n on γ_2 with $d_\Gamma(p_n, r_n) \leq \epsilon$. Then, since γ_2 is nf-straight, the subpath $\gamma_2^{(n)}$ of γ_2 from vertex r_{-n} to r_n is uniquely determined by r_{-n} and r_n for each $n \geq 0$. But the number of possibilities for r_{-n} and r_n for all $\gamma_2 \in \Pi$ is bounded above by the square K of the total number of words in Σ^* of length at most ϵ. So, since each γ_2 is the union of its subpaths $\gamma_2^{(n)}$, it follows that the number $|\Pi|$ of choices for γ_2 is bounded above by the same constant K.

To prove that Π is non-empty, we now construct a particular path γ_2 in Π via a sequence of vertices t_i on that path, using an argument attributed to Delzant. For each $n \geq 0$, we define ${\gamma'_1}^{(n)}$ to be the (unique) nf-reduced path joining p_{-n} to p_n. We know from the Bounded Coset Penetration Property that each such path ${\gamma'_1}^{(n)}$ lies in $\Delta(\gamma_1)$. Now, first we define the vertex t_0 to be a vertex with $d_\Gamma(t_0, p_0) \leq \epsilon$ that occurs in ${\gamma'_1}^{(n)}$ for infinitely many n. Then for each $m = 1, 2, \ldots$, we define the vertices t_m, t_{-m} together with a path $\gamma_2^{(m)}$ of Γ-length $2m$ joining t_{-m} to t_m such that (i) $d_\Gamma(t_{-m}, t_{-m+1}) = d_\Gamma(t_m, t_{m-1}) = 1$; (ii) $\gamma_2^{(m-1)}$ is a subpath of $\gamma_2^{(m)}$; and (iii) $\gamma_2^{(m)}$ is a subpath of ${\gamma'_1}^{(n)}$ for infinitely many n. Then we define the path γ_2 to be the union of the paths $\gamma_2^{(m)}$. Since

each path $\gamma_2^{(m)}$ is a subpath of ${\gamma_1'}^{(n)}$ for some n, we have $\gamma_2 \in \Delta(\gamma_1)$. To see that $\gamma_1 \in \Delta(\gamma_2)$, we observe that each non-internal vertex of γ_1 is at Δ-distance at most ϵ from a non-internal vertex in ${\gamma_1'}^{(n)}$ for all sufficiently large n and hence also from a non-internal vertex in $\gamma_2^{(m)}$ for all sufficiently large m. So γ_2 is in the set Π.

Now, fix some $\gamma_2 \in \Pi$, let r_0 be a vertex on γ_2 with $d_\Gamma(p_0, r_0) \leq \epsilon$, and let K be the upper bound on $|\Pi|$ found above. It is shown next in [7] that, for some M with $1 \leq M \leq K$, the isometry of Γ induced by multiplication on the left by u^M fixes γ_2. For suppose that I_u is the isometry of Γ induced by left multiplication within G by u. Since I_u fixes γ_1 it must fix $\Delta(\gamma_1)$ and so induce a permutation of the finite set Π. Hence for any γ_2 there is an M with $1 \leq M \leq K$ for which $I_u^M = I_{u^M}$ fixes γ_2.

Let $\alpha \in \Sigma^*$ label the path from p_0 to r_0, so $|\alpha| \leq \epsilon$. The isometry maps r_0 to the vertex of γ_2 labelled by $u^M\alpha$, which we denote by r_M. Since multiplication by u^M is an isometry of the labelled graph Γ, it must map the two vertices of γ_2 adjacent to r_0 to the two vertices adjacent to r_M. Since we know that the former are not labelled by generators in the same parabolic subgroup H_i, the same applies to the latter, so r_M also has the property of being a non-internal vertex of γ_2.

So the label z of the subpath of γ_2 from r_0 to r_M is nf-straight with $\alpha^{-1}u^M\alpha =_G z$. Since we can recognise nf-straight words in polynomial time, we can find suitable α and M in polynomial time by exhaustive search.

Algorithmic steps: find α and M as described, and compute an SLP for the nf-straight word z.

3.2.3 Testing conjugacy of the M th powers

The next step is to test for conjugacy between z and v^M. We follow the argument of [7, Section 3.3], and in some cases (when we found it appropriate) we use notation from that article.

Our Fig 2. is based on [7, Fig. 7]. In the diagram, n is some suitably large positive integer, and w and y are nf-reduced words in Σ^* labelling paths from the vertices e and p to the vertex q', respectively. Suppose that $\beta^{-1}z\beta =_G v^M$ (and hence $\beta^{-1}z^{2n}\beta =_G v^{2Mn}$ as shown in the diagram). We define q to be the vertex at the end of the path labelled v^{Mn} that starts at p; then q is the midpoint of the path labelled v^{2Mn} from p to q'.

As we remarked at the end of Section 3.2.1, the Bounded Coset Penetration Property implies the existence of a non-internal vertex r on the path labelled y that is at Γ-distance at most ϵ from the vertex q. Also, [12, Proposition 6.3] (which is the result that we need to apply the argument of [7, Section 3.3]

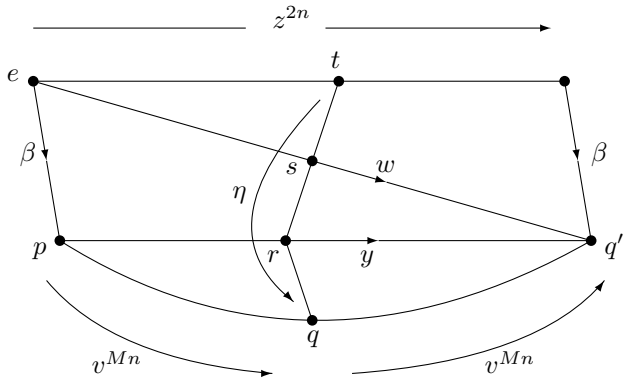

Figure 2: Testing conjugacy of M-th powers

in the context of relatively hyperbolic groups) tells us that, provided that n is sufficiently large compared with $|\hat{\beta}|$, there are non-internal vertices s and t on the paths labelled by the words w and z^{2n} respectively such that $d_\Gamma(r,s) + d_\Gamma(s,t)$ is at most a constant N (which is called L' in [12]). So the Γ-length of the word $\eta \in \Sigma^*$, labelling the path from t to vertex q via s and r, is bounded by the constant $N' := \epsilon + N$.

The group element labelling the path in the diagram from the vertex e to the vertex t is of the form $z^m z_1$ for some non-splitting prefix z_1 of z, and so $z^m z_1 \eta =_G \beta v^{Mn}$. Hence $\eta =_G z_1^{-1} z^{-m} \beta v^{Mn}$, and

$$\eta v^M \eta^{-1} =_G z_1^{-1} z^{-m} \beta v^M \beta^{-1} z^m z_1 =_G z_1^{-1} z z_1;$$

the right-hand-size freely reduces to a non-splitting cyclic conjugate of z.

It follows that to test for conjugacy between z and v^M we should check all possible conjugators η of Σ-length up to N', reducing each conjugate $\eta v^M \eta^{-1}$ to a word y in nf, and checking for each such y whether y is a non-splitting cyclic conjugate of z. As we explained in Section 2.4, we can decide in polynomial time (given SLPs for y, z) whether y is a non-splitting subword of z^2, and hence whether y is a non-splitting cyclic conjugate of z.

If z and v^M are not conjugate then neither are u and v, and so in that case the algorithm returns no. But otherwise we have found a conjugator η' (a product of η and a subword of z) such that $\eta' v^M \eta'^{-1} =_G u^M$, and we need to continue. In that case, u and v will be conjugate in G if and only if u and $\eta' v \eta'^{-1}$ are conjugate by an element of $C_G(u)$. So now we replace v by $\eta' v \eta'^{-1}$, and hence we may assume that $u^M =_G v^M =_G z$.

Algorithmic step: Check all possible conjugators of u^M to v_M as just described. If a conjugator is found, then store it (as SLP). If not then return no.

3.2.4 Completion of the proof

We have $u^M =_G v^M =_G z$, where as before z is the nf-straight word equal in G to u^M, and we want to decide whether u and v are conjugate in G. If so, then a conjugator will lie in the centraliser $C_G(z)$.

Now, we let β be a word over Σ representing an arbitrary element of C. Then $\beta^{-1}z\beta =_G v^M =_G z$, the diagram of Fig 2 applies also to these choices of u, v, z, β and M, and the argument of Section 3.2.3 applies just as before to give a non-splitting prefix z_1 of z, integers m, n, and a word η of length at most N', with $z^m z_1 \eta =_G \beta v^{Mn} =_G \beta z^n =_G z^n \beta$ (recall that $\beta \in C_G(z)$). It follows that $\beta =_G z^{m-n} z_1 \eta$.

Now we find the shortest non-splitting prefix z_0 of z that is a *root* of z: that is, for which there is an integer $\ell \geq 1$ so that $z = z_0^\ell$. We can find z_0 as the non-splitting prefix of z^2 that ends immediately before the second occurrence of z as a subword of the word z^2. Note that z_0 must be stable under cyclic derivation, because z is.

Then we define z_3 to be the minimal non-splitting suffix of z_1 for which $z_1 = z_0^{\ell'} z_3$ for some $\ell' \in \mathbb{Z}$. So $\beta =_G z_0^{\ell''} z_3 \eta$, with $\ell'' = \ell(m-n) + \ell' \in \mathbb{Z}$. We note that z_3 is also the unique shortest non-splitting prefix of z satisfying $\eta z \eta^{-1} =_G z_3^{-1} z z_3$; hence z_3 is completely determined by z and η. We observe that z_0 (as the root of z) is similarly uniquely determined, but z_1 is not necessarily. Hence once we have found η, we have determined the possible β.

In order to find β, we check all words of length at most N' over Σ as candidates for the associated word η, as follows. For each such (candidate) η, we compute $\mathsf{nf}(\eta z \eta^{-1})$ and test whether the result is a non-splitting cyclic conjugate of z. If so, then we define z_3 to be the shortest non-splitting prefix of z with $\eta z \eta^{-1} =_G z_3^{-1} z z_3$, and compute an SLP for $\mathsf{nf}(z_3 \cdot \eta)$ (which lies in $C_G(z)$); we store all such SLPs in a set C_z. Then $|C_z| \leq J$, where J is defined to be the number of words over Σ of length at most N'. Certainly the set C_z must contain the normal form of the product $z_3\eta$ of the previous paragraph. Hence β is equal in G to the product of a power of z_0 and an element z' of the set C_z.

We claim that we only need to check those elements $\beta =_G z_0^{\ell''} z'$ with $0 \leq \ell'' \leq (J-1)!$ and $z' \in C_z$ in order to locate an element $\beta \in C_G(z)$ that conjugates u to v (if such an element exists). This is because those elements between them cover the cosets in $C_G(z)$ of its centre, which has index at most $J!$ in $C_G(z)$, and all of whose elements centralise u and v.

If one of these elements conjugates u to v then we return yes and the combined conjugators from this and the previous sections, and otherwise we return no.

Algorithmic steps:

(i) Find the root z_0 of z as described above.

(ii) Construct the set C_z as described above.

(iii) For each $z' \in C_z$ and each ℓ'' with $0 \leq \ell'' \leq (J-1)!$, check whether $z_0^{\ell''} z'$ conjugates u to v. If so, then return yes and a conjugator.

(iv) If none of the elements tested in Step (iii) conjugates u to v then return no.

References

[1] J. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro and H. Short, "Notes on word-hyperbolic groups", in E. Ghys, A. Haefliger and A. Verjovsky, eds., *Proceedings of the Conference "Group Theory from a Geometric Viewpoint" held in I.C.T.P., Trieste, March 1990*, World Scientific, Singapore, 1991.

[2] Y. Antolín and L. Ciobanu, Finite generating sets of relatively hyperbolic groups and applications to geodesic languages. *Trans. Amer. Math. Soc.*, 368(11):7965–8010, 2016.

[3] B.H. Bowditch, Relatively hyperbolic groups, *Internat. J. Algebra Comput.*, 22(03):66pp, 2012.

[4] Martin R. Bridson and André Haefliger. *Metric Spaces of Non-Positive Curvature*. Springer Verlag, 1999.

[5] Moses Charikar, Eric Lehman, April Lehman, Ding Liu, Rina Panigrahy, Manoj Prabhakaran, Amit Sahai, and Abhi Shelat. The smallest grammar problem. *IEEE Transactions on Information Theory*, 51(7):2554–2576, 2005.

[6] D.B.A. Epstein, J. Cannon, D. Holt, D., S. Levy, M. Paterson, and W. Thurston. *Word Processing in Groups*, Jones and Bartlett, Boston, 1992.

[7] David Epstein and Derek Holt. The linearity of the conjugacy problem in word-hyperbolic groups. *Internat. J. Algebra Comput.*, 16(2):287–305, 2006.

[8] B. Farb, Relatively hyperbolic groups, *Geom. Funct. Anal.* 8(5):810–840, 1998.

[9] Christian Hagenah. *Gleichungen mit regulären Randbedingungen über freien Gruppen*. PhD thesis, University of Stuttgart, 2000.

[10] D. Holt, M. Lohrey, S. Schleimer, Compressed decision problems in hyperbolic groups, *arXiv:1808.06886*, Preprint, 2018.

[11] D. Holt, S. Rees, Regularity of quasigeodesics in a hyperbolic group, *Internat. J. Algebra Comput.* 13(5):585–596, 2003.

[12] D. Holt, S. Rees, The compressed word problem in relatively hyperbolic groups, submitted to *J. Algebra*, arXiv:2005.13917. *Internat. J. Algebra Comput.* 13(5):585–596, 2003.

[13] Marek Karpinski, Wojciech Rytter, and Ayumi Shinohara. Pattern-matching for strings with short descriptions. In *Combinatorial pattern matching (Espoo, 1995)*, volume 937 of *Lecture Notes in Comput. Sci.*, pages 205–214. Springer, Berlin, 1995.

[14] Markus Lohrey. Word problems and membership problems on compressed words. *SIAM Journal on Computing*, 35(5):1210 – 1240, 2006.

[15] Markus Lohrey. *The Compressed Word Problem for Groups*. SpringerBriefs in Mathematics. Springer, 2014.

[16] Denis V. Osin. Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems. *Mem. Amer. Math. Soc.*, 179(843):vi+100, 2006.

[17] Wojciech Plandowski. Testing equivalence of morphisms on context-free languages. In *Proceedings of ESA 1994*, volume 855 of *Lecture Notes in Computer Science*, pages 460–470. Springer, 1994.