Automorphic orbits in free groups: words versus subgroups

Pedro Silva, Pascal Weil

Automorphic orbits in free groups: words versus subgroups *

Pedro V. Silva, `pvsilva@fc.up.pt`
Centro de Matemática, Universidade do Porto[†]

Pascal Weil, `pascal.weil@labri.fr`
LaBRI, Université de Bordeaux and CNRS[‡]

31th March 2010

ABSTRACT

We show that the following problems are decidable in a rank 2 free group F_2: does a given finitely generated subgroup H contain primitive elements? and does H meet the orbit of a given word u under the action of G, the group of automorphisms of F_2? Moreover, decidability subsists if we allow H to be a rational subset of F_2, or alternatively if we restrict G to be a rational subset of the set of invertible substitutions (a.k.a. positive automorphisms). In higher rank, the following weaker problem is decidable: given a finitely generated subgroup H, a word u and an integer k, does H contain the image of u by some k-almost bounded automorphism? An automorphism is k-almost bounded if at most one of the letters has an image of length greater than k.

2000 Mathematics Subject Classification: 20E05

Orbit problems in general concern the orbit of an element u or a subgroup H of a group F, under the action of a subset G of $\mathsf{Aut}\,F$. Conjugacy problems are a special instance of such problems, where G consists of the inner automorphisms of F. In this paper, we restrict our attention to the case where F is the free group F_A with finite basis A.

In this context, orbit problems were maybe first considered by Whitehead [26], who proved that membership in the orbit of u under the action of $\mathsf{Aut}\,F_A$ is decidable. The analogous result regarding the orbit of a finitely generated subgroup H was established by Gersten [7]. Much literature has been devoted as well to the case where $G = \langle \varphi \rangle$ is a cyclic subgroup of $\mathsf{Aut}\,F_A$, e.g. Myasnikov and Shpilrain's work [15] on finite orbits of the form $\langle \varphi \rangle \cdot u$ and Brinkmann's recent proof [3] of the decidability of membership in $\langle \varphi \rangle \cdot u$.

The orbit problems considered in this paper are of the following form: given an element $u \in F_A$, a finitely generated subgroup H of F_A and a subset G of $\mathsf{Aut}\,F_A$, does H meet the orbit of u under the action of G; that is: does H contain $\varphi(u)$ for some automorphism $\varphi \in G$? A particular instance of this problem, when $G = \mathsf{Aut}\,F_A$, is the question whether H contains a primitive element, since the set of primitive elements of F_A is the automorphic orbit of each letter $a \in A$. The latter problem was recently solved by Clifford and Goldstein in full generality [4]. These problems were posed to the second author by O. Bogopolski, and they appear as Problem F39 in the list of open problems on `grouptheory.info`.

Our main results state that these problems are decidable in the rank 2 free group F_2, if $G = \mathsf{Aut}\,F_2$ (Theorem 2.3) or if G belongs to a certain family of rational subsets of $\mathsf{Aut}\,F_2$, which includes the rational subsets of invertible substitutions (a.k.a. positive automorphisms, which map each letter to a positive word) or of inverses of invertible substitutions, see Sections 5.1 and 5.2. For these rational values of G, we also show the decidability of subgroup orbit problems: given two finitely generated subgroups H, K of F_2, does there exist $\mu \in G$ such that K is contained in (resp. equal to) $\mu(H)$.

In free groups with larger rank, we are only able to decide a weaker problem. Say that an automorphism φ of F_A is k-almost bounded if $|\varphi(a)| > k$ for at most one letter $a \in A$. We show that given $k > 0$, $u \in F_A$ and H a finitely generated subgroup of F_A, one can decide whether there exists a k-almost bounded automorphism μ such that $\mu(u) \in H$.

Some of our results hold also if we replace the subgroup H by a rational subset of F_A.

We use two main methods. Some of our main results can be derived from

2

general results on the decidability of the solvability of equations with rational constraints in free groups (Diekert, Gutiérrez and Hagenah [6], building on Makanin's famous result [14]). This is an interesting application of equations in free groups.

For other results, we give a direct combinatorial proof. We use a particular factorization of the automorphism group $\mathsf{Aut}\,F_2$ (Theorem 3.4) and a detailed combinatorial analysis of the effect of certain simple automorphisms on the graphical representation of the subgroup H (the representation by means of so-called Stallings foldings [24, 10], see Section 1.2). The set of these automorphisms is $\Sigma = \{\varphi_{a,ba}, \varphi_{b^{-1},a^{-1}}, \varphi_{b,a}\}$ ($\varphi_{x,y}$ maps generator a to x and generator b to y).

This combinatorial analysis leads to the definition of a (large but finite) automaton whose vertices are finite automata associated with the Stallings automata of the subgroups in the Σ^*-orbit of H. The construction of this automaton exploits the fact that a certain combinatorial parameter of Stallings automata (which we call the number of singularities) is preserved under the action of automorphisms in Σ. And it is the possibility of reading these actions on this finite automaton which yields our decidability results for the cases where G is a rational subset of Σ^*. Invertible substitutions form a particular rational submonoid of Σ^*.

Interesting intermediary results state that the set of primitive elements in F_2 is a context-sensitive language (Proposition 3.8) and that if $|A| = m$ and $v_1, \ldots, v_{m-1} \in F_A$, then the set of elements x such that $v_1, \ldots, v_{m-1}, x$ form a basis of F_A is a constructible rational set (Proposition 2.11).

1   Preliminaries

1.1   Free groups

Let A denote a finite alphabet. The *free monoid on A*, written A^*, is the set of all finite sequences of elements of A (including the empty sequence, written 1), under the operation of concatenation. We also write A^+ for the set of non-empty sequences of elements of A.

Let A^{-1} be a disjoint set of formal inverses of A and let $\tilde{A} = A \cup A^{-1}$. The operation $u \mapsto u^{-1}$ is extended to $\tilde{A}^*$ as usual, by letting $(a^{-1})^{-1} = a$ and $(ua)^{-1} = a^{-1}u^{-1}$ for all $a \in A$ and $u \in \tilde{A}^*$.

The *free group* on A is the quotient F_A of $\tilde{A}^*$ by the congruence generated by the pairs $(aa^{-1}, 1)$, $a \in \tilde{A}$, and we write $\pi\colon \tilde{A}^* \to F_A$ for the canonical projection. A word $u \in \tilde{A}^*$ is *reduced* if it does not contain a factor aa^{-1} ($a \in \tilde{A}$) and we denote by R_A the set of reduced words. We also say that

$u \in R_A$ is *cyclically reduced* if uu is reduced as well. And we denote by $cc(u)$ the *cyclic core* of u, that is, the unique word such that u is of the form $u = v^{-1}cc(u)v$ in $\tilde{A}^*$.

We write $u \mapsto \bar{u}$ the *reduction map*, where $\bar{u}$ is the (uniquely defined) word obtained from u by iteratively deleting factors of the form aa^{-1} ($a \in \tilde{A}$) until none is left. It is well-known that the reduction map is well defined, and that the restriction $\pi \colon R_A \to F_A$ is a bijection. To simplify notation, if $g \in F_A$, we also write $\bar{g}$ for the reduced word such that $\pi(\bar{g}) = g$, and we let the *length* of g be $|g| = |\bar{g}|$.

Given $X \subseteq F_A$, we denote by $\langle X \rangle$ the subgroup of F_A generated by X. We also let $\mathsf{Aut}\, F_A$ denote the automorphism group of F_A. If $\varphi \in \mathsf{Aut}\, F_A$ and no confusion arises, we shall denote also by φ the corresponding bijection of R_A.

Given $B \subseteq F_A$, we say that B is a *basis* of F_A if the homomorphism from F_B to F_A induced by the inclusion map $B \to F_A$ is an isomorphism. Equivalently, B is a basis of F_A if and only if $B = \varphi(A)$ for some $\varphi \in \mathsf{Aut}\, F_A$. The *primitive* elements of F_A are those that sit in some basis of F_A. It follows that the set of primitive elements of F_A is the orbit of each letter $a \in A$ under the action of $\mathsf{Aut}\, F_A$.

In much of this paper, we shall be discussing the free group on 2 generators. We fix the alphabet $A_2 = \{a, b\}$ and use the notation $F_2 = F_{A_2}$, $R_2 = R_{A_2}$.

1.2 Automata and rational subsets

The product of two subsets K, L of a monoid M is the subset $KL = \{xy \mid x \in K,\ y \in L\}$. The *star* operator on subsets is defined by $L^* = \bigcup_{n \geq 0} L^n$, where $L^0 = \{1\}$. A subset L of a monoid is said to be *rational* if L can be obtained from finite subsets using finitely many times the operators union, product and star. We denote by $\mathsf{Rat}\, M$ the set of rational subsets of M. If M is the free monoid A^* on a finite alphabet A, subsets of A^* are called *languages*, and elements of $\mathsf{Rat}\, A^*$ are called *rational languages*.

Note that if $\varphi \colon A^* \to M$ is an onto morphism, then $\mathsf{Rat}\, M$ is the set of all $\varphi(L)$ where $L \in \mathsf{Rat}\, A^*$. For instance, every finitely generated subgroup of F_A is rational.

It is well-known that rational languages can be characterized by means of finite automata. A (finite) A-*automaton* is a tuple $\mathcal{A} = (Q, q_0, T, E)$ where Q is a (finite) set, $q_0 \in Q$, $T \subseteq Q$ and $E \subseteq Q \times A \times Q$. It can be viewed as a graph with vertex set Q (the *states*), with a designated vertex q_0 (the *initial state*) and a set of designated vertices T (the *terminal* states), whose edges

are labeled by letters in A, and are given by the set E (the *transitions*).

A *nontrivial path* in $\mathcal{A}$ is a sequence

$$p_0 \xrightarrow{a_1} p_1 \xrightarrow{a_2} \dots \xrightarrow{a_n} p_n$$

with $n \geq 1$, $(p_{i-1}, a_i, p_i) \in E$ for $i = 1, \dots, n$. Its *label* is the word $a_1 \dots a_n \in A^+$. We consider also the *trivial path* $p_0 \xrightarrow{1} p_0$ for each $p_0 \in Q$, whose label is the empty word. A path is said to be *successful* if $p_0 = q_0$ and $p_n \in T$. The *language $L(\mathcal{A})$ recognized by $\mathcal{A}$* is the set of all labels of successful paths in $\mathcal{A}$.

The automaton $\mathcal{A} = (Q, q_0, T, E)$ is said to be *deterministic* if, for all $p \in Q$ and $a \in A$, there is at most one edge of the form (p, a, q). In that case, we write $q = p \cdot a$. We say that $\mathcal{A}$ is *trim* if every $q \in Q$ lies in some successful path.

Kleene's theorem states that a language is rational if and only if it is accepted by a finite automaton, which can be required to be deterministic and trim, see [9]. In the context of a particular result or claim, we say that a rational language L is *effectively constructible* if there exists an algorithm to produce a finite automaton recognizing L from the concrete structures containing the input. More generally, if $\varphi \colon A^* \to M$ is an onto morphism, we say that a rational subset of M is *effectively constructible* (with respect to A) if it is the image of an effectively constructible rational language over A.

Remark 1.1 A subset $L \subseteq F_A$ is rational if $L = \pi(K)$ for some rational subset K of $\tilde{A}^*$. Benois' theorem [1] states that this is the case if and only if $\overline{K}$ $(= \pi^{-1}(\pi(K)) \cap R_A$, in bijection with L via $\pi)$ is a rational subset of $\tilde{A}^*$. In the sequel we sometimes confuse the notions of a rational subset of F_A and a rational language in $\tilde{A}^*$ that consists only of reduced words.

1.3 Automata and subgroups of F_A

To discuss subgroups of free groups, we use inverse automata. In an $\tilde{A}$-automaton $\mathcal{A} = (Q, q_0, T, E)$, the *dual* of an edge $(p, a, q) \in E$ is (q, a^{-1}, p). Then $\mathcal{A}$ is said to be *dual* if E contains the duals of all edges, and *inverse* if it is dual, deterministic, trim (equivalent to *connected* in this case) and $|T| = 1$.

Given a finitely generated subgroup H of F_A (we write $H \leq_{\mathsf{fg}} F_A$), we denote by $\mathcal{A}(H)$ the *Stallings automaton* associated to H by the construction often referred to as *Stallings foldings*. This construction, that can be traced
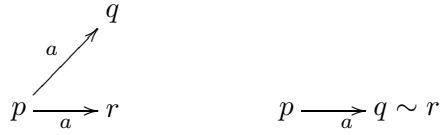
Figure 1: A folding step, with $a \in \tilde{A}$

back to the early part of the twentieth century [19, Chapter 11], was made explicit by Serre [20] and Stallings [24] (see also [10]).

A brief description is as follows. If $h_1, \ldots, h_r \in R_A$ is a set of generators of the subgroup H (that is, $H = \langle \pi(h_1), \ldots, \pi(h_r) \rangle$), one constructs a dual automaton in the form of r subdivided circles around a common distinguished vertex 1 (both initial and terminal), each labeled by one of the h_i. Then we iteratively identify identically labeled pairs of edges starting (resp. ending) at the same vertex (this is called the *folding* process, see Figure 1), until no further folding is possible.

The following proposition summarizes important properties, see [10].

Proposition 1.2 *Let $H \leq_{fg} F_A$. Then:*

(i) *$\mathcal{A}(H)$ is a finite inverse automaton, which does not depend on the finite reduced generating set nor on the sequence of foldings chosen;*

(ii) *if $p \xrightarrow{u} q$ is a path in $\mathcal{A}(H)$, so is $p \xrightarrow{\overline{u}} q$;*

(iii) *for every $u \in R_A$, $u \in L(\mathcal{A}(H))$ if and only if $\pi(u) \in H$; in particular, $L(\mathcal{A}(H)) \subseteq \pi^{-1}(H)$;*

(iv) *for every cyclically reduced $u \in F_A$, $wuw^{-1} \in H$ for some $w \in F_A$ if and only if u labels some loop in $\mathcal{A}(H)$.*

2 The mixed orbit problem as an equation problem

Our original motivation on writing this paper was solving the mixed orbit problem $\varphi(u) \in H$ for given $u \in F$ and $H \leq_{fg} F$. We managed to solve it in rank 2, see Section 2.1 below. We also solve it in arbitrary rank, if we impose a restriction on the class of automorphisms, see Section 2.2 below.

Our initial proof in rank 2 was purely combinatorial, and is presented in Section 5.3 below, whereas our result in higher ranks made use of a

6

result of Diekert, Gutiérrez and Hagenah [6] on equations with rational constraints which we will discuss below. Upon reading a first version of this paper[1], Dahmani and Girardel, and independently Enric Ventura (recalling a foregone conversation with Alexei Miasnikov), called our attention to the fact that Diekert, Gutiérrez and Hagenah's result could also be used to prove our result in rank 2.

2.1 Equations with rational constraints and automorphisms of F_2

A *system of equations* in a free group F_A, with set of unknowns X (disjoint from A), is a tuple $\mathcal{E} = (e_1, \ldots, e_k)$ of elements of $F_{A \cup X}$. A *solution* of $\mathcal{E}$ is a morphism from $F_{A \cup X}$ to F_A which maps each letter of A to itself and every element of $\mathcal{E}$ to 1.

A *rational constraint* on a system of equations with unknowns in X is a collection $\mathcal{L} = (L_x)_{x \in X}$ of rational subsets of F_A and we say that a morphism $\varphi \colon F_{A \cup X} \to F_A$ is a solution of the system $\mathcal{E}$ with rational constraints $\mathcal{L}$ if φ is a solution of $\mathcal{E}$ and $\varphi(x) \in L_x$ for each $x \in X$. Diekert, Gutiérrez and Hagenah [6] showed the following result.

Theorem 2.1 *The satisfiability problem for systems of equations with rational constraints in a free group is decidable.*

This leads to a quick solution of the mixed orbit problem in F_2.

Theorem 2.2 *Given $u \in F_2$ and a rational subset L of F_2, it is decidable whether or not $\varphi(u) \in L$ for some $\varphi \in \mathsf{Aut}\, F_2$.*

Proof. By a result attributed to Dehn, Magnus and Nielsen (see [22]), $\{x, y\}$ is a basis of F_2 if and only if there exists some $g \in F_2$ such that $g^{-1}[x, y]g = [a, b]^{\pm 1}$.

Now observe that if $u \in F_2$, then $\varphi(u) \in H$ for some $\varphi \in \mathsf{Aut}\, F_2$ if and only if $u(x, y) \in H$ for some basis $\{x, y\}$ — where $u(x, y)$ denotes the word u in which each occurrence of a has been replaced by x and each occurrence of b by y.

Thus there exists an automorphism φ such that $\varphi(u) \in H$ if and only if one of the following systems (in the unknowns x, y, v, g, and for $\varepsilon = \pm 1$) admits a solution

$$\begin{cases} g^{-1}[x, y]g = [a, b]^\varepsilon \\ u(x, y) = v \end{cases}$$

with the rational constraint that $v \in L$. This is decidable by Theorem 2.1.
$\square$

Since a finitely generated subgroup is rational (it is the star of its generators and their inverses), Theorem 2.2 yields the following corollary.

Corollary 2.3 *Given $u \in F_2$ and $H \leq_{\mathsf{fg}} F_2$, it is decidable whether or not $\varphi(u) \in H$ for some $\varphi \in \mathsf{Aut}\, F_2$.*

Since the primitive elements of F_2 are the orbit of each letter $a \in A$ under $\mathsf{Aut}\, F_2$, we also note the following result.

Corollary 2.4 *Given a rational subset L of F_2 (e.g. a finitely generated subgroup), it is decidable whether or not L contains a primitive element.*

Remark 2.5 Clifford and Goldstein also proved a comparable result for primitive elements, by completely different methods: they show that it is decidable whether a finitely generated subgroup H of F_A (for any finite alphabet A) contains a primitive element [4].

One can also consider, in the statement of Theorem 2.2, the rational subset of positive elements of F_2 (namely, the submonoid A^*). Then our result shows that it is decidable whether an element $u \in F_2$ is *potentially positive*, that is, whether it has a positive automorphic image. Different proofs of this result already appear in Goldstein [8] and Lee [11].

Another idea is to consider a tuple of elements of F_2 rather than a single element u, or equivalently a subgroup $K \leq_{\mathsf{fg}} F_2$.

Theorem 2.6 *Let $u_1, \ldots, u_k \in F_2$, $L_1, \ldots, L_k \in \mathsf{Rat}\, F_2$ and $H, K \leq_{\mathsf{fg}} F_2$. The following problems are decidable:*

(1) whether $\varphi(u_1) \in L_1, \ldots, \varphi(u_k) \in L_k$, for some $\varphi \in \mathsf{Aut}\, F_2$;

(2) whether conjugates of $\varphi(u_1), \ldots, \varphi(u_k)$ sit in $L_1, \ldots, L_k$, respectively, for some $\varphi \in \mathsf{Aut}\, F_2$;

(3) whether $\varphi(K) \subseteq H$, for some $\varphi \in \mathsf{Aut}\, F_2$.

Proof. These statements are proved like Theorem 2.2, by reduction to a system of equations with rational constraints.

For (1), we consider the system

$$
\begin{cases}
g^{-1}[x,y]g = [a,b]^\varepsilon \\
u_1(x,y) = v_1 \\
\quad \dots \\
u_k(x,y) = v_k
\end{cases}
$$

in the unknowns $x, y, g, v_1, \dots, v_k$ with the rational constraints $v_1 \in L_1, \dots, v_k \in L_k$.

For (2), we consider the system

$$
\begin{cases}
g^{-1}[x,y]g = [a,b]^\varepsilon \\
h_1^{-1}u_1(x,y)h_1 = v_1 \\
\quad \dots \\
h_k^{-1}u_k(x,y)h_k = v_k
\end{cases}
$$

in the unknowns $x, y, g, h_1, \dots, h_k, v_1, \dots, v_k$ with the rational constraints $v_1 \in L_1, \dots, v_k \in L_k$.

Statement (3) is a particular case of (1), when $u_1, \dots, u_k$ is a generating set of K and $L_1 = \dots = L_k = H$. $\square$

Remark 2.7 Instead of asking whether there exists an automorphism in $\mathsf{Aut}\, F_2$ mapping u into L, one may want to exhibit such an automorphism, if it exists.

The existence question reduces to the satisfiability of equations with rational constraints, and Diekert, Gutiérrez and Hagenah showed that this can be done in PSPACE [6]. One can extract from that paper a description of such a solution (an automorphism) as an exponential length product of simple automorphisms: the images of the letters may therefore have double exponential length.

2.2 Beyond rank 2

We do not know how to extend Theorem 2.2 or Corollary 2.3 to arbitrary finite alphabets, but we can get decidability for weakened versions of the problem. The first such result involves a restriction on the subgroups considered.

Theorem 2.8 *Let $u \in F_A$ and let $H \leq_{fg} F_A$. If H is cyclic or a free factor of F_A, it is decidable whether or not $\varphi(u) \in H$ for some $\varphi \in \mathsf{Aut}\, F_A$.*

Proof. Let us first assume that H is a free factor of F_A, with rank k. It is easily verified that $\varphi(u) \in H$ for some automorphism φ if and only if u sits in some rank k free factor of F_A. This is known to be decidable: it suffices to compute a minimum length element v in the automorphic orbit of u (the so-called easy part of Whitehead's algorithm, see [13, 18]) and to verify whether v uses at least k letters (Shenitzer [21], see also [13, Prop. I.5.4]).

Let us now assume that $H = \langle v \rangle$. Without loss of generality, we may assume that u and v are cyclically reduced. Say that a word x is root-free if it is not equal to a non-trivial power of a shorter word. Then $u = x^k$ for some uniquely determined integer $k \geq 1$ and root-free word x, and similarly, $v = y^\ell$ for some uniquely determined $\ell \geq 1$ and root-free y. It is an elementary verification that the image of a root-free word by an automorphism is also root-free. Thus, an automorphism maps u into H if and only if it maps x to y or y^{-1}, and k is a multiple of ℓ. Decidability follows from the fact that we can decide whether two given words are in each other's automorphic orbit, using Whitehead's algorithm [13]. $\square$

The second result on a weakened version of our orbit problem involves *almost bounded automorphisms*. Given a finite alphabet A and $k \in \mathbb{N}$, we say that an automorphism φ of F_A is k-almost bounded if $|\varphi(a)| > k$ for at most one letter $a \in A$. We let $\mathsf{AlmB}_k F_A$ denote the set of k-almost bounded automorphisms of F_A.

Theorem 2.9 *Given $u \in F_A$, $L \in \mathsf{Rat}\, F_A$ and $k \in \mathbb{N}$, it is decidable whether or not $\varphi(u) \in L$ for some $\varphi \in \mathsf{AlmB}_k F_A$.*

The proof of this theorem relies on Diekert, Gutiérrez and Hagenah's result on the satisfiability of equations with rational constraints in free groups discussed in Section 2.1. We also require two technical results.

Lemma 2.10 *Let $A = \{a_1, \ldots, a_m\}$ and $u \in R_A$. Then $\{a_1, \ldots, a_{m-1}, u\}$ is a basis of F_A if and only if $u = v a_m^\varepsilon w$ for some $v, w \in R_{\{a_1, \ldots, a_{m-1}\}}$ and $\varepsilon \in \{1, -1\}$.*

Proof. It is immediate that if $u = v a_m^\varepsilon w$ with $v, w \in R_{\{a_1, \ldots, a_{m-1}\}}$, then $\{a_1, ..., a_{m-1}, u\}$ generates F_A, and by the Hopfian property of free groups (see [13, Prop. I.3.5]), $\{a_1, ..., a_{m-1}, u\}$ is a basis of F_A.

Conversely, let $u \in R_A$ contain at least an occurrence of a_m or a_m^{-1}, and let $u = vzw$ be the factorization with $v, w \in R_{\{a_1, \ldots, a_{m-1}\}}$ of maximal length. It is immediate that if $H = \langle a_1, ..., a_{m-1}, u \rangle$, then $H = \langle a_1, ..., a_{m-1}, z \rangle$ and

$\mathcal{A}(H)$ is equal to $\mathcal{A}(\langle z \rangle)$ with loops labelled $a_1, ..., a_{m-1}$ attached at the origin. Thus, if $\{a_1, ..., a_{m-1}, u\}$ is a basis of F_A, then $\mathcal{A}(\langle z \rangle)$ must consist of a single loop labeled a_m, and hence z must be equal to a_m or a_m^{-1}. $\square$

This leads to the following generalization.

Proposition 2.11 *Let* $m = |A|$ *and* $v_1, \ldots, v_{m-1} \in R_A$. *Then*

$$X = \{x \in R_A \mid (v_1, \ldots, v_{m-1}, x) \text{ is a basis of } F_A\}$$

is rational and effectively constructible.

Proof. First note that X is nonempty if and only if $(v_1, ..., v_{m-1})$ is a basis of a free factor of F_A. This is well-known to be decidable. Moreover, if $X \neq \emptyset$, then we can effectively construct an element z of X: it is verified in [23] that if $K = \langle v_1, ..., v_{m-1} \rangle$, then K is a free factor of F_A if and only if there are vertices p and q of $\mathcal{A}(K)$ whose identification leads (via foldings) to the bouquet of circles $\mathcal{A}(F_A)$, and in that case, if u_p and u_q are the labels of geodesic paths of $\mathcal{A}(K)$ from the origin to p and q, then $z = \overline{u_p u_q^{-1}} \in X$.

Let $\varphi \in \text{Aut}\, F_A$ be defined by $\varphi(a_i) = v_i$ $(i = 1, \ldots, m-1)$ and $\varphi(a_m) = z$. Then $x \in X$ if and only if $(a_1, \ldots, a_{m-1}, \varphi^{-1}(x))$ is a basis of F_A. By Lemma 2.10, this is equivalent to say that $\varphi^{-1}(x) \in R(a_m \cup a_m^{-1})R$, where $R = \langle a_1, \ldots, a_{m-1} \rangle$, and therefore

$$X = \varphi(R(a_m \cup a_m^{-1})R) = V(z \cup z^{-1})V$$

for $V = \langle v_1, \ldots, v_{m-1} \rangle$.

In particular, X is rational and the formula $X = V(z \cup z^{-1})V$ provides an effective construction for it. $\square$

Proof of Theorem 2.9. Write $A = \{a_1, \ldots, a_m\}$. Without loss of generality, we may restrict ourselves to the case $|\varphi(a_i)| \leq k$ for $i = 1, \ldots, m-1$. Since there are only finitely many choices for these $\varphi(a_i)$, we may as well assume them to be fixed, say $\varphi(a_i) = v_i$ for $i = 1, \ldots, m-1$. Let then $X = \{x \in R_A \mid (v_1, \ldots, v_{m-1}, x) \text{ is a basis of } F_A\}$: then X is rational by Proposition 2.11.

Write $u = u_0 a_m^{\varepsilon_1} u_1 \ldots a_m^{\varepsilon_n} u_n$ with $n \geq 0$, $u_i \in F_{\{a_1, \ldots, a_{m-1}\}}$ and $\varepsilon_i = \pm 1$ for every i. Then we must decide whether there exists some $y \in X$ such that

$$u_0' y^{\varepsilon_1} u_1' \ldots y^{\varepsilon_n} u_n' \in L,$$

where $u_i' = u_i(v_1, \ldots, v_{m-1})$ is the word obtained from u_i by replacing each a_j by v_j. This is equivalent to deciding whether the equation

$$u_0' y^{\varepsilon_1} u_1' \ldots y^{\varepsilon_n} u_n' = z \tag{1}$$

on the variables y, z has a solution in F_A with the rational constraints $y \in X$ and $z \in L$. This is decidable by Theorem 2.1. $\square$

As in Section 2.1, Theorem 2.9 yields a decidability result for finitely generated subgroups.

Corollary 2.12 *Given $u \in F_A$, $H \leq_{fg} F_A$ and $k \in \mathbb{N}$, it is decidable whether or not $\varphi(u) \in H$ for some $\varphi \in AlmB_k \, F_A$.*

And as in Theorem 2.6, we use the same ideas to prove the following theorem. If $w \in F_2$, λ_w denotes the inner automorphism $u \mapsto w^{-1}uw$.

Theorem 2.13 *Let $u_1, \ldots, u_m \in F_A$, $L \in Rat \, F_A$, $H, K \leq_{fg} F_A$ and $k \in \mathbb{N}$. The following problems are decidable:*

(1) whether $\varphi(u_1), \ldots, \varphi(u_m) \in L$, for some $\varphi \in AlmB_k \, F_A$;

(2) whether $\lambda_w \varphi(u_1), \ldots, \lambda_w \varphi(u_m) \in L$, for some $w \in F_A$ and $\varphi \in AlmB_k \, F_A$;

(3) whether conjugates of $\varphi(u_1), \ldots, \varphi(u_m)$ sit in L, for some $\varphi \in AlmB_k \, F_A$;

(4) whether $\varphi(K) \subseteq H$, for some $\varphi \in AlmB_k \, F_A$;

(5) whether $\lambda_w \varphi(K) \subseteq H$, for some $w \in F_A$ and $\varphi \in AlmB_k \, F_A$.

Proof. These statements are proved like Theorem 2.9, by reduction to a system of equations with rational constraints.

For the first statement, we consider a system of equations of the form of equation (1) in the proof of Theorem 2.9, one for which u_j, $1 \leq j \leq m$ (the unknowns are $y, z_1, \ldots, z_m$).

For the second (resp. third) statement, we consider the same system, with each equation conjugated by a new unknown v (resp. by distinct new unknowns v_j, $1 \leq j \leq m$).

The fourth and fifth statements are applications of the first and second when the rational subset L is the subgroup H. $\square$

Remark 2.14 Following-up with the discussion in Remark 2.7, we note that the complexity upper bounds for the decision problems described in this section are PSPACE again: we need to (attempt to) solve, successively, systems of equations for the different values of $v_1, \ldots, v_{m-1}$ (with the notation of the proof of Theorem 2.9. These $(m-1)$-tuples of words of length at most k are exponentially many (in the variable k) but they can be listed in polynomial space.

3 Combinatorial approach: the role of Σ

We now restrict our attention to F_2. If $x, y \in F_2$, we denote by $\varphi_{x,y}$ the endomorphism mapping a to x and b to y. In this section, we discuss some properties of the following sets of automorphisms of F_2:

$$\Sigma_0 = \left\{ \varphi_{a,ba}, \varphi_{b^{-1},a^{-1}} \right\} \text{ and } \Sigma = \Sigma_0 \cup \left\{ \varphi_{b,a} \right\};$$

$$\Phi = \left\{ \varphi_{a,ba}, \varphi_{ab,b}, \varphi_{a,ab}, \varphi_{ba,b} \right\};$$

$$\Delta = \left\{ \varphi_{a,a^m b^\varepsilon a^n} \mid m, n \in \mathbb{Z}, \ \varepsilon \in \{1, -1\} \right\};$$

$$\Psi = \left\{ \varphi \in \mathsf{Aut}\, F_2 : |\varphi(a)| = |\varphi(b)| = 1 \right\} \text{ and } \Lambda = \left\{ \lambda_w \mid w \in R_A \right\}.$$

The following will be useful in the sequel.

Proposition 3.1 *(i) $X\Lambda = \Lambda X$ for every $X \subseteq \mathsf{Aut}\, F_2$;*

(ii) $\Lambda \Psi \Phi^* \subseteq \Lambda \Psi (\Sigma_0^{-1})^* \varphi_{a^{-1},b}$;

(iii) $\Delta \subseteq \Lambda (\varphi_{a,ba}^* \cup \varphi_{a^{-1},b} \varphi_{a,ba}^* \varphi_{a^{-1},b})(1 \cup \varphi_{a,b^{-1}})$.

Proof. (i) follows from the fact that $\theta \lambda_w = \lambda_{\theta(w)} \theta$ for each $w \in F_2$ and $\theta \in \mathsf{Aut}\, F_2$.

(ii) Notice that $\varphi_{ab,b} = \varphi_{b,a} \varphi_{a,ba} \varphi_{b,a}$, $\varphi_{a,ab} = \lambda_{a^{-1}} \varphi_{a,ba}$ and $\varphi_{ba,b} = \lambda_{b^{-1}} \varphi_{ab,b}$. It follows that $\Lambda \Psi \Phi^* \subseteq \Lambda \Psi \{\varphi_{a,ba}, \varphi_{b,a}\}^*$.

Observe also that $\varphi_{a,ba} = \varphi_{a^{-1},b} \varphi_{a,ba}^{-1} \varphi_{a^{-1},b}$, $\varphi_{b,a} = \varphi_{a^{-1},b} \varphi_{b^{-1},a^{-1}}^{-1} \varphi_{a^{-1},b}$ and $\varphi_{a^{-1},b}^2 = 1$. So we have

$$\{\varphi_{a,ba}, \varphi_{b,a}\}^* = \varphi_{a^{-1},b} \{\varphi_{a,ba}^{-1}, \varphi_{b^{-1},a^{-1}}^{-1}\}^* \varphi_{a^{-1},b} = \varphi_{a^{-1},b} (\Sigma_0^{-1})^* \varphi_{a^{-1},b}.$$

Therefore $\Lambda \Psi \Phi^* \subseteq \Lambda \Psi \varphi_{a^{-1},b} (\Sigma_0^{-1})^* \varphi_{a^{-1},b} = \Lambda \Psi (\Sigma_0^{-1})^* \varphi_{a^{-1},b}.$

13

(iii) Observe that if $m, n \in \mathbb{Z}$, then $\varphi_{a,a^m b a^n} = \lambda_{a^{-m}} \varphi_{a,ba^{m+n}} = \lambda_{a^{-m}} \varphi_{a,ba}^{m+n}$, so that $\varphi_{a,a^m b a^n} \in \Lambda(\varphi_{a,ba}^* \cup (\varphi_{a,ba}^{-1})^*)$. We already noted that $\varphi_{a,ba}^{-1} = \varphi_{a^{-1},b} \varphi_{a,ba} \varphi_{a^{-1},b}$ and $\varphi_{a^{-1},b}^2 = 1$, so

$$\varphi_{a,a^m b a^n} \in \Lambda(\varphi_{a,ba}^* \cup \varphi_{a^{-1},b} \varphi_{a,ba}^* \varphi_{a^{-1},b}).$$

Similarly, $\varphi_{a,a^m b^{-1} a^n} = \lambda_{a^n} \varphi_{a,ba}^{-(m+n)} \varphi_{a,b^{-1}}$ and hence

$$\varphi_{a,a^m b^{-1} a^n} \in \Lambda(\varphi_{a,ba}^* \cup \varphi_{a^{-1},b} \varphi_{a,ba}^* \varphi_{a^{-1},b}) \varphi_{a,b^{-1}},$$

which concludes the proof. $\square$

3.1 Primitive words and a factorization of $\mathsf{Aut}\, F_2$

Let us first consider a particular automorphic orbit in F_A, namely the set P_A of primitive words. Recall that a word is *primitive* if it belongs to some basis of F_A. In particular, P_A is the automorphic orbit of each letter from A. We shall often view P_A as a subset of R_A. We denote by P_2 the set of all primitive words in F_2.

We use a known characterization of the words in P_2 to derive a technical factorization of the group $\mathsf{Aut}\, F_2$ of automorphisms of F_2, that will be used in Section 5. We further exploit this characterization to point out certain language-theoretic properties of P_2.

Proposition 3.2 reports two results: the first is due to Nielsen [16] (see also [5, 2.2] and [17]) and the second is due to Wen and Wen [25]. An interesting perspective on either is offered in [12, Chapter 2] and [2, Chapter I-5].

Proposition 3.2 *(i) Up to conjugation, every primitive element $u \in P_2$ is either a letter, or of the form $u = a^{n_1} b^{m_1} ... a^{n_k} b^{m_k}$ where*

- *either $n_1 = ... = n_k \in \{1, -1\}$ and $\{m_1, ..., m_k\} \subseteq \{n, n+1\}$ for some integer n,*
- *or $m_1 = ... = m_k \in \{1, -1\}$ and $\{n_1, ..., n_k\} \subseteq \{n, n+1\}$ for some integer n.*

(ii) The set of positive primitive words $P_2 \cap \{a, b\}^+$ is equal to $\Phi^*(\{a, b\}) = b \cup \Phi^*(a)$.

Corollary 3.3 $P_2 = \Lambda\Psi\Phi^*(a)$.

Proof. By Proposition 3.2 (i), every primitive element of F_2 is a conjugate of $\psi(ab^{m_1}...ab^{m_k})$, where $\{m_1, ..., m_k\} \subseteq \{n, n+1\}$ for some integer $n \geq 0$ and $\psi \in \Psi$. That is, $P_2 = \Lambda\Psi(P_2 \cap \{a, b\}^+)$. By Proposition 3.2 (ii), it follows that $P_2 = \Lambda\Psi(b \cup \Phi^*(a)) = \Lambda\Psi\Phi^*(a)$. $\square$

We can now prove a useful decomposition result for $\mathsf{Aut}\, F_2$.

Theorem 3.4 $Aut\, F_2 = \Lambda\Psi\Phi^*\Delta = \Psi(\Sigma_0^{-1})^*\Lambda\varphi_{a,ba}^*(\varphi_{a^{-1},b} \cup \varphi_{a^{-1},b^{-1}})$.

Proof. To establish the first equality, we consider $\theta \in \mathsf{Aut}\, F_2$. Then $\theta(a) \in P_2$ and so $\theta(a) = \sigma(a)$ for some $\sigma \in \Lambda\Psi\Phi^*$ by Corollary 3.3. Corollary 2.10 then shows that $\sigma^{-1}\theta = \varphi_{a,a^m b^\varepsilon a^n}$ for some $m, n \in \mathbb{Z}$ and $\varepsilon \in \{1, -1\}$. So $\sigma^{-1}\theta \in \Delta$ and $\theta \in \Lambda\Psi\Phi^*\Delta$. It follows that

$$
\begin{aligned}
\mathsf{Aut}\, F_2 &\subseteq \Lambda\Psi\Phi^*(\varphi_{a,ba}^* \cup \varphi_{a^{-1},b}\varphi_{a,ba}^*\varphi_{a^{-1},b})(1 \cup \varphi_{a,b^{-1}}) \text{ by Proposition 3.1} \\
&\subseteq \Lambda\Psi\Phi^*(1 \cup \varphi_{a^{-1},b}\varphi_{a,ba}^*\varphi_{a^{-1},b})(1 \cup \varphi_{a,b^{-1}}) \text{ since } \varphi_{a,ba} \in \Phi \\
&\subseteq \Lambda\Psi\Phi^*(\varphi_{a^{-1},b}\varphi_{a,ba}^*\varphi_{a^{-1},b})(1 \cup \varphi_{a,b^{-1}}) \text{ since } \varphi_{a^{-1},b}^2 = 1 \\
&\subseteq \Lambda\Psi(\Sigma_0^{-1})^*\varphi_{a,ba}^*\varphi_{a^{-1},b}(1 \cup \varphi_{a,b^{-1}}) \text{ by Proposition 3.1 (ii)} \\
&\subseteq \Psi(\Sigma_0^{-1})^*\Lambda\varphi_{a,ba}^*(\varphi_{a^{-1},b} \cup \varphi_{a^{-1},b^{-1}}).
\end{aligned}
$$

The converse inclusion is of course trivial. $\square$

3.2 Invertible substitutions

A *substitution*[2] of F_A is an endomorphism φ such that $\varphi(a) \in A^*$ for every $a \in A$. If φ is an automorphism, it is said to be an *invertible substitution*. We denote by $\mathrm{IS}(F_2)$ the monoid of all invertible substitutions of F_2, and by $\mathrm{IS}^{-1}(F_2)$ the monoid of their inverses. Note that the inverse of an invertible substitution is not necessarily a substitution: indeed $\varphi_{a,ba}^{-1} = \varphi_{a,ba^{-1}}$.

Lemma 3.5 $IS(F_2)$ *is a rational submonoid of* Σ^*. *Moreover, there exists a rational submonoid* S *of* Σ^* *such that* $IS^{-1}(F_2) = \varphi_{a,b^{-1}} S \varphi_{a,b^{-1}}$.

In addition, every rational subset $R \in \mathsf{Rat}\, IS(F_2)$ *is also in* $\mathsf{Rat}\, \Sigma^*$, *and every rational subset* $R \in \mathsf{Rat}\, IS^{-1}(F_2)$ *is of the form* $\varphi_{a,b^{-1}} R' \varphi_{a,b^{-1}}$ *for some* $R' \in \mathsf{Rat}\, S \subseteq \mathsf{Rat}\, \Sigma^*$.

---

[2] also called a *positive endomorphism*

Proof. It is known [25] that the monoid $\mathrm{IS}(F_2)$ is generated by $\varphi_{b,a}$, $\varphi_{a,ba}$ and $\varphi_{a,ab}$ (see also [2, Chapter I.5], [12, Sec. 2.3.5]). But $\varphi_{b,a}, \varphi_{a,ba} \in \Sigma$ and

$$\varphi_{a,ab} = \varphi_{b,a}\varphi_{b^{-1},a^{-1}}\varphi_{a,ba}\varphi_{b^{-1},a^{-1}}\varphi_{b,a} \in \Sigma^*,$$

so $\mathrm{IS}(F_2) = \{\varphi_{b,a}, \varphi_{a,ba}, \varphi_{a,ab}\}^* \in \mathsf{Rat}\,\Sigma^*$.

Next we observe that

$$\varphi_{b,a}^{-1} = \varphi_{b,a} = \varphi_{a,b^{-1}}\varphi_{b^{-1},a^{-1}}\varphi_{a,b^{-1}}$$
$$\varphi_{a,ba}^{-1} = \varphi_{a,ba^{-1}} = \varphi_{a,b^{-1}}\varphi_{a,ab}\varphi_{a,b^{-1}}$$
$$\varphi_{a,ab}^{-1} = \varphi_{a,a^{-1}b} = \varphi_{a,b^{-1}}\varphi_{a,ba}\varphi_{a,b^{-1}}$$

Since $\varphi_{a,b^{-1}}$ has order 2, it follows that $\mathrm{IS}(F_2)^{-1} = \varphi_{a,b^{-1}}R\varphi_{a,b^{-1}}$ with $R = \{\varphi_{b^{-1},a^{-1}}, \varphi_{a,ab}, \varphi_{a,ba}\}^* \in \mathsf{Rat}\,\Sigma^*$. $\square$

3.3 Primitive words form a context-sensitive language

Digressing from our main topic, we use Corollary 3.3 to establish a language-theoretic property of primitive words.

Recall that a *context-sensitive A-grammar* is a triple $\mathcal{G} = (V, P, S)$ where V is a finite set containing A, S is an element of V that is not in A and P is the *set of rules* of the grammar: a finite set of pairs $(\ell, r) \in V^+ \times V^+$ such that

$$\ell \notin A^+ \text{ and } |\ell| \le |r|.$$

For all $x, y \in V^+$, we write $x \Rightarrow y$ if there exist $u, v \in V^*$ and $(\ell, r) \in P$ such that $x = u\ell v$ and $y = urv$. We denote by $\overset{*}{\Rightarrow}$ the transitive and reflexive closure of $\Rightarrow$. The language *generated by $\mathcal{G}$* is

$$L(\mathcal{G}) = \{w \in A^+ \mid S \overset{*}{\Rightarrow} w\}.$$

A language $L \subseteq A^+$ is said to be *context-sensitive* if it is generated by some context-sensitive A-grammar. As usual, a language $L \subseteq A^*$ is called *context-sensitive* if $L \cap A^+$ is context-sensitive.

The right and left quotients of a language L by a word u are defined by

$$u\backslash L = \{x \in A^* \mid ux \in L\}, \quad L/u = \{x \in A^* \mid xu \in L\}.$$

Lemma 3.6 *The class of context-sensitive languages is closed under union, intersection, concatenation, right and left quotient by a word, 1-free substitutions and inverse morphisms.*

Proof. Closure under union, intersection, concatenation, 1-free substitutions, and inverse homomorphisms is well-known [9, Exercise 9.10]. In particular, the family of context-sensitive languages forms a *trio* [9, Section 11.1] and as such, it is closed under *limited erasing* [9, Lemma 11.2]. By definition, this means that if $k \geq 1$, L is context-sensitive and φ is a morphism such that $\varphi(v) \neq 1$ for each $u \in L$ and each factor v of u of length greater than k, then $\varphi(L)$ is context-sensitive as well.

For the quotients, it suffices to consider letters, hence let $L \subseteq A^*$, $a \in A$ and $\$ \notin A$. Let σ be the substitution that maps a to $\sigma(a) = \{a, \$\}$ and which fixes every other letter of A. Let also $\varphi \colon (A \cup \{\$\})^* \to A^*$ be the morphism which fixes every letter of A and erases $\$$. Then $a \backslash L = \varphi(\sigma(L) \cap \$ A^*)$ and $L/a = \varphi(\sigma(L) \cap A^* \$)$. Since the σ-images of the letters are finite, and hence context-sensitive, the languages $\sigma(L) \cap \$ A^*$ and $\sigma(L) \cap A^* \$$ are context-sensitive; moreover φ exhibits limited erasing on these languages, so $a \backslash L$ and L/a are context-sensitive as well. $\square$

Proposition 3.7 *Let A be a finite alphabet and let Γ be a finite set of endomorphisms of A^+. For every $u \in A^+$, $\Gamma^*(u)$ is a context-sensitive language.*

Proof. Take $b \notin A$. We define a context-sensitive $(A \cup \{b\})$-grammar $\mathcal{G} = (V, P, S)$ by $V = A \cup \{R, S, T\} \cup \{F_\varphi \mid \varphi \in \Gamma\}$ and

$$P = \{S \to bF_\varphi uR, \ S \to bub^2, \ F_\varphi a \to \varphi(a)F_\varphi, \ F_\varphi R \to TR,$$
$$F_\varphi R \to b^2, \ aT \to Ta, \ bT \to bF_\varphi; \ a \in A, \ \varphi \in \Gamma \ \}.$$

We show that $L(\mathcal{G}) = b\Gamma^*(u)b^2$.

Clearly, $F_\varphi v \overset{*}{\Rightarrow} \varphi(v)F_\varphi$ for all $\varphi \in \Gamma$ and $v \in A^*$ and so

$$bvTR \overset{*}{\Rightarrow} bTvR \Rightarrow bF_\varphi vR \overset{*}{\Rightarrow} b\varphi(v)F_\varphi R \Rightarrow b\varphi(v)TR.$$

Since $S \Rightarrow bF_\varphi uR \overset{*}{\Rightarrow} b\varphi(u)F_\varphi R \Rightarrow b\varphi(u)TR$ for every $\varphi \in \Gamma$, it follows that $S \overset{*}{\Rightarrow} b\theta(u)F_\varphi R \Rightarrow b\theta(u)b^2$ for every $\theta \in \Gamma^+$. Together with $S \Rightarrow bub^2$, this yields $b\Gamma^*(u)b^2 \subseteq L(\mathcal{G})$.

To prove the opposite inclusion, let

$$Z = \big\{S\big\} \cup \big\{bxyb^2, bxTyR, b\varphi(x)F_\varphi yR \mid xy \in \Gamma^*(u)\big\}.$$

Then Z is closed under $\Rightarrow$. That is: if $X \in Z$ and $X \Rightarrow Y$, then $Y \in Z$.

Since $S \in Z$, it follows that $L(\mathcal{G}) \subseteq Z \cap A^* = b\Gamma^*(u)b^2$ and so $L(\mathcal{G}) = b\Gamma^*(u)b^2$. Thus $b\Gamma^*(u)b^2$ is context-sensitive and by Lemma 3.6, $\Gamma^*(u) = b\backslash(b\Gamma^*(u)b^2)/b^2$ is context-sensitive as well. $\square$

Theorem 3.8 $\overline{P_2}$ *is a context-sensitive language.*

Proof. Since the class of context-sensitive languages is closed under union (Lemma 3.6), it follows from Proposition 3.2(ii) and Proposition 3.7 that $P_2 \cap \{a, b\}^+ = \overline{P_2} \cap \{a, b\}^+$ is context-sensitive. Moreover, Proposition 3.2(i) shows that $P_2 = \Lambda\Psi(P_2 \cap \{a, b\}^+) = \Psi\Lambda(P_2 \cap \{a, b\}^+)$. Since Ψ is finite, we need only prove that each $\overline{\psi\Lambda(P_2 \cap \{a, b\}^+)}$, $\psi \in \Psi$, is context-sensitive.

Notice that, for each $\psi \in \Psi$ and each word w, $\overline{\psi(w)} = \psi(\overline{w})$. By Lemma 3.6 again, we need only to prove that $\overline{\Lambda(P_2 \cap \{a, b\}^+)}$ is context-sensitive.

Let $w \in R_2$ and $p \in P_2 \cap \{a, b\}^+$. If wpw^{-1} is not reduced, then one of wp and pw^{-1} is not reduced. In the first case, let q be the longest prefix of p such that q^{-1} is a suffix of w, say $p = qr$ and $w = vq^{-1}$. Then $wpw^{-1} = vq^{-1}qrqv^{-1} = vrqv^{-1}$. The second case (if wp is reduced but pw^{-1} is not) is treated similarly. Iterating this reasoning, we find that $wpw^{-1} = vp'v^{-1}$, where v is a prefix of w and p' is a cyclic shift of the word p – that is, there are words q, r such that $p = qr$ and $p' = rq$.

Since $P_2 \cap \{a, b\}^+$ is closed under taking cyclic shifts, it follows that $\Lambda(P_2 \cap \{a, b\}^+)$ is the set of reduced words of the form vpv^{-1} with $p \in P_2 \cap \{a, b\}^+$.

Thus, if $\mathcal{G} = (V, P, S)$ is a context-sensitive A-grammar generating $P_2 \cap \{a, b\}^+$, then $\overline{\Lambda(P_2 \cap \{a, b\}^+)} = L(\mathcal{G}') \cap R_2$, where $\mathcal{G}' = (V', P', S')$ is the context-sensitive A-grammar given by $S' \notin V$, $V' = \{S'\} \cup V$ and $P' = P \cup \{S' \to S\} \cup \{S' \to cS'c^{-1}; \ c \in A_2 \cup A_2^{-1}\}$. In view of the closure properties in Lemma 3.6, $\overline{\Lambda(P_2 \cap \{a, b\}^+)}$ is context-sensitive, and hence so is P_2 . $\square$

This result cannot be improved to the next level of Chomsky's hierarchy:

Proposition 3.9 $\overline{P_2}$ *is not a context-free language.*

Proof. We show that $P_2 \cap ab^+ab^+ab^+$ is not a context-free language. Since the class of context-free languages is closed under intersection with rational languages, it shows that P_2 is not context-free either.

It follows easily from Proposition 3.2(i) that $P_2 \cap ab^*ab^*ab^*$ is equal to

$$\left\{ ab^m ab^n ab^k \mid m, n, k \in \mathbb{N}, \ \max(m, n, k) = \min(m, n, k) + 1 \right\}. \quad (2)$$

It is now a classical exercise to show that $P_2 \cap ab^+ab^+ab^+$ is not context-free since it fails the Pumping Lemma for context-free languages [9, Section 6.1]. $\square$

4 Singularities, bridges and automorphisms in Σ

We now discuss the evolution of the Stallings automaton of a subgroup H under the iterated action of the automorphisms in Σ. It is well-known that the automata $\mathcal{A}(\varphi(H))$ may grow unboundedly as the length of φ (as a product of elements of Σ) grows. But in the context of the mixed orbit problem with respect to the automorphisms in Σ^*, we are only interested in the possibility of reading a u-labeled loop (where u is a fixed word) in $\mathcal{A}(\varphi(H))$: if the growth of the automata results in long stretches without branchpoints, then this growth does not affect the membership of u in $\varphi(H)$ after a certain point.

Indeed, we show that the fragments of the $\mathcal{A}(\varphi(H))$ ($\varphi \in \Sigma^*$) that could conceivably allow the reading of a u-loop take only finitely many values – and these fragments (which we call truncated automata) can be organised as the states of an automaton on alphabet Σ. The mixed orbit problem with respect to Σ^* then reduces to deciding whether this automaton accepts a non-empty language.

We now get into the technical considerations that give substance to this overview of our method. Given $H \leq_{\mathsf{fg}} F_2$, we say that a state q of $\mathcal{A}(H)$ is

- a *source* if $q \cdot a, q \cdot b \neq \emptyset$, $\qquad\qquad \overset{a}{\leftarrow} q \overset{b}{\longrightarrow}$

- a *sink* if $q \cdot a^{-1}, q \cdot b^{-1} \neq \emptyset$. $\qquad \overset{a}{\longrightarrow} q \overset{b}{\leftarrow}$

Note that a source may have incoming edges and a sink may have outgoing edges. We use the general term *singularities* to refer to both sources and sinks and we denote by $\mathrm{Sing}(H)$ the set of all singularities of $\mathcal{A}(H)$ plus the origin.

If we emphasize the vertices of $\mathrm{Sing}(H)$ in $\mathcal{A}(H)$, it is immediate that $\mathcal{A}(H)$ can be described as the union of *positive paths*, i.e. paths with label in $(a \cup b)^+$, between the vertices of $\mathrm{Sing}(H)$, and these positive paths do not intersect each other except at $\mathrm{Sing}(H)$. We call such paths *bridges*. Note that every positive path whose internal states are not singularities can be extended into a uniquely determined bridge.

4.1 Bridges in $\mathcal{A}(H)$

The next two results are easily verified.

Fact 4.1 *The automaton $\mathcal{A}(\varphi_{b^{-1},a^{-1}}(H))$ has the same vertex set as $\mathcal{A}(H)$, edges are reverted and labels changed. In particular, sources and sinks are exchanged. If β is a bridge in $\mathcal{A}(H)$, $\beta = p \overset{w}{\longrightarrow} q$, then there is a bridge*

of equal length $q \longrightarrow p$ in $\mathcal{A}(\varphi_{b^{-1},a^{-1}}(H))$, labeled $\varphi_{b^{-1},a^{-1}}(w^{-1})$, which we denote by $\varphi_{b^{-1},a^{-1}}(\beta)$.

Fact 4.2 *The automaton $\mathcal{A}(\varphi_{b,a}(H))$ has the same vertex set as $\mathcal{A}(H)$ and labels are exchanged. Sources and sinks remain the same. If β is a bridge in $\mathcal{A}(H)$, $\beta = p\xrightarrow{w}q$, then there is a bridge of equal length $p \longrightarrow q$ in $\mathcal{A}(\varphi_{b,a}(H))$, labeled $\varphi_{b,a}(w)$, which we denote by $\varphi_{b,a}(\beta)$.*

Dealing with $\varphi_{a,ba}$ is naturally a little more complex. However, as we will see in the next two statements, the foldings implied in computing $\mathcal{A}(\varphi_{a,ba}(H))$ are very local: they can be performed in a single round of independent foldings. Moreover, sources in $\mathcal{A}(\varphi_{a,ba}(H))$ were already sources in $\mathcal{A}(H)$ and sinks in $\mathcal{A}(\varphi_{a,ba}(H))$ are at distance 1 of sinks in $\mathcal{A}(H)$.

Fact 4.3 *The automaton $\mathcal{A}(\varphi_{a,ba}(H))$ is obtained from $\mathcal{A}(H)$ by the following 3 steps:*

(S1) *If $p\xrightarrow{b}q$ is an edge of $\mathcal{A}(H)$ and q is not a sink, we replace that edge by a path $p\xrightarrow{b}\bullet\xrightarrow{a}q$, adding a new intermediate vertex for each such edge.*

(S2) *If $p\xrightarrow{b}q\xleftarrow{a}r$ is a sink in $\mathcal{A}(H)$, we replace this configuration by*

$$p \xrightarrow{\makebox[2cm]{b}} q \underset{a}{\rightleftharpoons} r$$
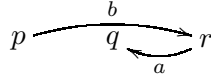
(S3) *We iteratively remove all the vertices of degree 1 different from the origin.*

Proof. Following [18, Subsection 1.2], the automaton $\mathcal{A}(\varphi(H))$ may be obtained from $\mathcal{A}(H)$ in three steps:

(1) We replace each edge labelled by b by a path labelled ba (introducing a new intermediate vertex for each such edge), producing a dual automaton $\mathcal{B}$.

(2) We execute the complete folding of $\mathcal{B}$.

(3) We successively remove all the vertices of degree 1 different from the origin.

How much folding is involved in the process? Let us consider the first level of folding, i.e. those pairs of edges that can be immediately identified in $\mathcal{B}$.

- There are no b-edges involved in the first level of folding: indeed, the b-edges keep their origin when we go from $\mathcal{A}(H)$ to $\mathcal{B}$, and their target is always a new vertex where folding cannot take place.

- If we have a sink $p\overset{b}{\longrightarrow}q\overset{a}{\longleftarrow}r$ in $\mathcal{A}(H)$, we get

$$p\overset{b}{\longrightarrow} \bullet \overset{a}{\longrightarrow}q\overset{a}{\longleftarrow}r$$

  in $\mathcal{B}$ and therefore an instance of first level folding, yielding

$$p \overset{b}{\underset{a}{\rightleftarrows}} q \longrightarrow r$$

- These are the only instances of first level folding: we cannot fold two "new" a-edges $\overset{a}{\longrightarrow}q\overset{a}{\longleftarrow}$ in $\mathcal{B}$ since that would imply the existence of two b-edges $\overset{b}{\longrightarrow}q\overset{b}{\longleftarrow}$ in $\mathcal{A}(H)$.

Let $\mathcal{C}$ denote the automaton obtained by performing all the instances of first level folding in $\mathcal{B}$. It follows from the above remarks that $\mathcal{C}$ can be obtained from $\mathcal{A}(H)$ by application of (S1) and (S2).

We actually need no second level of folding because $\mathcal{C}$ is already deterministic. Indeed, it is clear from (S1) and (S2) that configurations such as $\overset{a}{\longleftarrow}q\overset{a}{\longrightarrow}$ or $\overset{b}{\longleftarrow}q\overset{b}{\longrightarrow}$ cannot occur in $\mathcal{C}$.

Suppose that $\overset{b}{\longrightarrow}q\overset{b}{\longleftarrow}$ does occur. Then both edges must have been obtained through (S2) and the origin of these edges is the vertex $q \cdot (ab^{-1})$ in $\mathcal{A}(H)$, a contradiction.

Finally, suppose that $\overset{a}{\longrightarrow}q\overset{a}{\longleftarrow}$ does occur. At least one of these edges must have been obtained through (S1), but not both, otherwise we would have a configuration $\overset{b}{\longrightarrow}q\overset{b}{\longleftarrow}$ in $\mathcal{A}(H)$. But then we would have a configuration $\overset{a}{\longrightarrow}q\overset{b}{\longleftarrow}$ in $\mathcal{A}(H)$ and q would be a sink, contradicting the application of (S1). Thus $\mathcal{C}$ is deterministic and so $\mathcal{A}(\varphi(H))$ is obtained from $\mathcal{A}(H)$ by successive application of (S1), (S2) and (S3). $\square$

Fact 4.4 *(i) When applying $\varphi_{a,ba}$, a state of $\mathcal{A}(H)$ is trimmed in step (S3) if and only if it is a sink of $\mathcal{A}(H)$ without outgoing edges. Moreover, no consecutive states can be trimmed.*

(ii) The sources of $\mathcal{A}(\varphi_{a,ba}(H))$ are precisely the sources p of $\mathcal{A}(H)$ such that $p \cdot a$ is not a sink or has outgoing edges in $\mathcal{A}(H)$.

(iii) The sinks of $\mathcal{A}(\varphi_{a,ba}(H))$ are precisely the states p of $\mathcal{A}(H)$ with incoming edges such that $p \cdot a$ is a sink of $\mathcal{A}(H)$.

Proof. (i) The origin cannot be trimmed and the number of outgoing edges never decreases, so the only possible candidates to (S3) are the states that see a decrease in their number of incoming edges, which are precisely the sinks of $\mathcal{A}(H)$. Their fate will then depend on the previous existence of some outgoing edge. Note that $\mathcal{A}(H)$ cannot possess two consecutive sinks with no outgoing edges, hence the trimming of a vertex will not be followed by the trimming of any of its neighbours.

(ii) Since outgoing edges can be at most redirected through (S1) and (S2), it is clear that every source p of $\mathcal{A}(\varphi_{a,ba}(H))$ must be a source of $\mathcal{A}(H)$. Thus everything will depend on $p \cdot a$ being trimmed or not, and part (i) yields the claim.

(iii) No new intermediate vertex obtained through (S1) can become a sink, and any sink of $\mathcal{A}(H)$ will not remain such after application of (S2). Thus the only remaining candidates are the non-sinks of $\mathcal{A}(H)$ that see an increase of their number of incoming edges, which are precisely those of the form $q \cdot a^{-1}$, where q is a sink of $\mathcal{A}(H)$. Clearly, to have two distinct incoming edges in $\mathcal{A}(\varphi_{a,ba}(H))$, $p = q \cdot a^{-1}$ must have at least one incoming edge in $\mathcal{A}(H)$. In such a case, it is easy to check that after (S1)/(S2), p has indeed become a sink of $\mathcal{A}(\varphi_{a,ba}(H))$. We remark also that the subsequent trimming by (S3) does not affect the presence of singularities. $\square$

Fact 4.5 *Let $\beta = p \overset{w}{\longrightarrow} q$ be a bridge in $\mathcal{A}(H)$ of length at least 2, and let $w = w'cd$ where $c, d \in A$.*

(i) *$\mathcal{A}(\varphi_{a,ba}(H))$ has a positive path $p \overset{\varphi_{a,ba}(w'c)}{\longrightarrow} s$, which extends to a uniquely determined bridge, denoted by $\varphi_{a,ba}(\beta)$.*

(ii) *$|\varphi_{a,ba}(\beta)| \geq |\beta| - 1$, and we have $|\varphi_{a,ba}(\beta)| = |\beta| - 1$ exactly if $w \in a^+$, p is a source or the origin in $\mathcal{A}(H)$, and q is a sink in $\mathcal{A}(H)$.*

Proof. Write $\beta = p \overset{w'}{\longrightarrow} r \overset{c}{\longrightarrow} s \overset{d}{\longrightarrow} q$.

(i) By Fact 4.4, no state of the path $p \overset{\varphi_{a,ba}(w'c)}{\longrightarrow} s$ risks trimming. Hence it suffices to check that no internal state of this path can become a singularity. This follows easily from Fact 4.4 (ii) and (iii).

(ii) The inequality $|\varphi_{a,ba}(\beta)| \geq |\beta| - 1$ follows at once from part (i). It follows also that $|\varphi_{a,ba}(\beta)| = |\beta| - 1$ if and only if $w'c \in a^+$ (otherwise $|\varphi_{a,ba}(\beta)| \geq |\varphi_{a,ba}(w'c)| > |w'c| = |\beta| - 1$) and $p, s \in \mathrm{Sing}(\varphi_{a,ba}(H))$. Thus we assume that $w'c \in a^+$.

Clearly, if p is the origin, it must remain so. If p is a source, it follows from Fact 4.4 (ii) that p remains a source (since $p \cdot a$ is not a sink in $\mathcal{A}(H)$).

Finally, if p is a sink, it will no longer be a singularity in $\mathcal{A}(\varphi_{a,ba}(H))$ by Fact 4.4 (iii). Therefore $p \in \mathrm{Sing}(\varphi_{a,ba}(H))$ if and only if it is a source or the origin in $\mathcal{A}(H)$.

Similarly, q can never become the origin or a source. Since q has incoming edges in $\mathcal{A}(H)$, it follows from Fact 4.4(iii) that s becomes a sink in $\mathcal{A}(\varphi_{a,ba}(H))$ if and only if $s \cdot a$ is a sink in $\mathcal{A}(H)$. Since the unique outgoing edge of s in $\mathcal{A}(H)$ has label d, then $s \in \mathrm{Sing}(\varphi_{a,ba}(H))$ if and only if $d = a$ and q is a sink in $\mathcal{A}(H)$. $\square$

4.2   Homogeneous cycles and cycle-free paths

Let $\sigma(H) = \max(1, source(H) + sink(H))$, where $source(H)$ (resp. $sink(H)$) is the number of sources (resp. sinks) of $\mathcal{A}(H)$. We call $\sigma(H)$ the *number of singularities* of $\mathcal{A}(H)$. Note that a vertex may be a source and a sink, and in that case, it contributes twice to $\sigma(H)$.

We say that a path $p \xrightarrow{w} r$ is *homogeneous* if $w \in R_a \cup R_b$, and it is *special homogeneous* if, in addition, it starts at a source or the origin, and it ends at a sink or the origin. Let $hc(\mathcal{A})$ (resp. $hcfp(\mathcal{A})$, $shcfp(\mathcal{A})$) be the maximum length of a homogeneous cycle (resp. homogeneous cycle-free path, special homogeneous cycle-free path) in automaton $\mathcal{A}$.

Given $H \leq_{f.g.} F_2$, we define

$$\delta_0(H) = \max(\sigma(H), hc(\mathcal{A}(H)),$$
$$\delta(H) = \max(\delta_0(H), hcfp(\mathcal{A}(H))),$$
$$\zeta(H) = \max(\delta_0(H), shcfp(\mathcal{A}(H))).$$

We record the following inequalities.

Lemma 4.6 *Let $H \leq_{f.g.} F_2$. Every cycle or a cycle-free path labeled b^k in $\mathcal{A}(\varphi_{a,ba}(H))$ satisfies $k \leq \sigma(H)$.*

Proof. Let us first assume that $\alpha = p \xrightarrow{b^k} q$ is a cycle-free path, say

$$p = q_0 \xrightarrow{b} q_1 \xrightarrow{b} \ldots \xrightarrow{b} q_k = q.$$

Since any b-edge obtained through (S1) must be followed only by an a-edge (see Fact 4.3), only the last edge $q_{k-1} \xrightarrow{b} q_k$ may be obtained through (S1), and the other edges arise from applications of (S2). Thus there exist edges

in $\mathcal{A}(H)$ (represented through discontinuous lines) of the form

$$
\begin{array}{ccccccc}
& p_1 & & p_2 & & p_{k-2} & & p_{k-1} \\
& \nearrow \uparrow & & \nearrow \uparrow & & \uparrow & \nearrow \uparrow \\
b \nearrow & |a & b \nearrow & |a & & |a & b \nearrow & |a \\
\nearrow & | & \nearrow & | & & | & \nearrow & | \\
q_0 \xrightarrow{\;b\;} q_1 \xrightarrow{\;b\;} q_2 & & \cdots & & q_{n-2} \xrightarrow{\;b\;} q_{k-1} \xrightarrow{\;b\;} q_k
\end{array}
$$

In particular, the vertices $p_1, \dots, p_{k-1}$ are distinct sinks in $\mathcal{A}(H)$, and the vertices $q_1, \dots, q_{k-1}$ are distinct sources in $\mathcal{A}(H)$. Therefore $2k - 2 \leq \sigma(H)$ and hence $k \leq \sigma(H)$.

If α is a cycle, then not even the last edge of α arises from an application of (S1), and the same reasoning shows that $2k \leq \sigma(H)$, so $k \leq \sigma(H)$. $\square$

Lemma 4.7 *Let $H \leq_{f.g.} F_2$ and $\varphi \in \Sigma$. Then*

$$
\sigma(\varphi(H)) \leq \sigma(H),
$$
$$
\delta_0(\varphi(H)) \leq \delta_0(H),
$$
$$
\zeta(\varphi(H)) \leq \zeta(H).
$$

Proof. The first inequality is a direct consequence of Facts 4.1, 4.2 and 4.4.

By Facts 4.1 and 4.2, the other inequalities are trivial if $\varphi = \varphi_{b,a}$ or $\varphi_{b^{-1},a^{-1}}$. We now assume that $\varphi = \varphi_{a,ba}$. Since $\sigma(\varphi(H)) \leq \sigma(H)$, we only need to show that the maximum length of a homogeneous cycle (resp. cycle-free special homogeneous) path $\alpha = p \longrightarrow q$ in $\mathcal{A}(\varphi(H))$ ($p = q$ in the case of a cycle) is at most equal to $\delta_0(H)$ (resp. $\zeta(H)$).

If the label of α is b^k, then Lemma 4.6 shows that $k \leq \sigma(H)$, so $k \leq \delta_0(H) \leq \zeta(H)$.

Suppose now that the label of α is a^k. In view of Fact 4.3, none of its edges was obtained trough (S1): indeed the a-edge in $\bullet \xrightarrow{\;b\;} \bullet \xrightarrow{\;a\;} \bullet$ produced by (S1) cannot occur in a homogeneous cycle, nor in a homogeneous path unless it is its first edge. But its initial vertex is not a singularity, so this edge cannot occur in a special homogeneous path. Hence the path α already existed in $\mathcal{A}(H)$. If α is a cycle, then $k \leq \delta_0(H)$.

If instead α is a special homogeneous cycle-free path, then Fact 4.4 (ii) shows that p is either the origin or a source in $\mathcal{A}(H)$. If q is the origin, we immediately get $k \leq \zeta(H)$. If instead q is a sink in $\mathcal{A}(\varphi(H))$, then $s = q \cdot a$ is a sink of $\mathcal{A}(H)$ by Fact 4.4 (iii), and we have a path

$$
\alpha' \;=\; p \xrightarrow{\;a^k\;} q \xrightarrow{\;a\;} s
$$

in $\mathcal{A}(H)$. If α' is cycle-free, then $k < k + 1 \leq \zeta(H)$. If, on the contrary, α' is not cycle-free, then s is the only repetition since the length k prefix of α', namely α, is cycle-free. If $s \neq p$, then q would also be a repetition since $\mathcal{A}(H)$ is an inverse automaton. Therefore $s = p$, so α' is a homogeneous cycle in $\mathcal{A}(H)$ and hence $k < k + 1 \leq \delta_0(H) \leq \zeta(H)$. This concludes the proof. $\square$

Remark 4.8 Note that it is not the case that $\delta(\varphi(H)) \leq \delta(H)$ always holds when $\varphi \in \Sigma$: see the case where $H = \langle ba \rangle$ and $\varphi = \varphi_{a,ba}$.

4.3 Truncated automata

Given $H \leq_{f.g.} F_2$, we consider the *geodesic metric d* defined on the vertex set of $\mathcal{A}(H)$ by taking $d(u, v)$ to be the length of the shortest path connecting u and v. Since $\mathcal{A}(H)$ is inverse, it is irrelevant to consider directed or undirected paths. As usual, we have

$$d(u, \operatorname{Sing}(H)) = \min\{d(u, v) \mid v \in \operatorname{Sing}(H)\}.$$

Given $t > 0$, the *t-truncation* of $\mathcal{A}(H)$, denoted by $\mathcal{A}_t(H)$, is the automaton obtained by removing from $\mathcal{A}(H)$ all vertices u such that $d(u, \operatorname{Sing}(H)) > t$ and their adjacent edges. Note that this automaton does not need to be connected.

We first observe that if β is a bridge which is long enough to be affected by the t-truncation of $\mathcal{A}(H)$, then for each $\varphi \in \Sigma$, $\varphi(\beta)$ is affected by the t-truncation of $\mathcal{A}(\varphi(H))$ as well.

Proposition 4.9 *Let $\varphi \in \Sigma$, $H \leq_{f.g.} F_2$ and $K \in \Sigma^*(H)$. If β is a bridge in $\mathcal{A}(K)$ and $|\beta| > \zeta(H)$, then $|\varphi(\beta)| \geq |\beta|$.*

Proof. The result is trivial if $\varphi = \varphi_{b^{-1}, a^{-1}}$ or $\varphi = \varphi_{b,a}$ since in those cases, $|\varphi(\beta)| = |\beta|$ (Facts 4.1 and 4.2). We now assume that $\varphi = \varphi_{a,ba}$.

By Fact 4.5, if $|\varphi(\beta)| < |\beta|$, then $\beta = p \xrightarrow{a^k} q$, where $k > \zeta(H)$, p is a source or the origin in $\mathcal{A}(K)$, and q is a sink of $\mathcal{A}(K)$. In particular, β is a special homogeneous cycle-free path, so that $|\beta| \leq \zeta(K)$.

Since $K \in \Sigma^* H$, Lemma 4.7 shows that $\zeta(K) \leq \zeta(H)$, a contradiction. $\square$

Theorem 4.10 *Let $\varphi \in \Sigma$, $H \leq_{f.g.} F_2$, $t > \frac{1}{2}\zeta(H)$ and $K, K' \in \Sigma^*(H)$.*
Then

$$\mathcal{A}_t(K) = \mathcal{A}_t(K') \implies \mathcal{A}_t(\varphi(K)) = \mathcal{A}_t(\varphi(K')).$$

Proof. As in several previous proofs, the result is trivial if $\varphi = \varphi_{b,a}$ or $\varphi_{b^{-1},a^{-1}}$, and we may assume that $\varphi = \varphi_{a,ba}$.

By Proposition 4.9, we know that, once the length of a bridge reaches the threshold $\zeta(H)+1$, it can only get longer. Since $t > \frac{1}{2}\zeta(H)$, t-truncation affects only bridges of length at least $\zeta(H) + 1$. We must therefore discuss the truncation mechanism for such long bridges.

Assume that $\beta = p \xrightarrow{w} q$ is a bridge in $\mathcal{A}(\mu(H))$ ($\mu \in \Sigma^*$) with $|w| \geq 2t + 1$. Then we may write $w = uzv$ with $|u| = |v| = t$. By Proposition 4.9, the label of $\varphi(\beta)$ is of the form $u'z'v'$ with $|u'| = |v'| = t$ and $|z'| \geq |z|$. We only need to prove that u' and v' depend only on $\mathcal{A}_t(\mu(H))$ and are therefore independent from z.

In view of Fact 4.4, it is clear that u' depends only on $\mathcal{A}_t(\mu(H))$ (remember that $w = uzv$ is a positive word and singularities cannot *move forward* along a positive path). The nontrivial case is of course the case of q being a sink in $\mathcal{A}(\mu(H))$, since by Fact 4.4 (iii) a sink can actually be transferred to the preceding state along a positive path. We claim that even in this case v' is independent from z.

Indeed, assume first that b occurs in v. Then $|\varphi(v)| > |v|$ provides enough compensation for the sink moving backwards one position. Hence we may assume that $v = a^t$. We claim that $v' = a^t$ as well, independently from z. Suppose not. Since we are assuming that the sink has moved from q to its predecessor, and $\varphi(a^{t-1}) = a^{t-1}$, it follows that $v' = ba^{t-1}$. Hence b occurs in w. Write $w = xba^m$. Since $\varphi(ba^m) = ba^{m+1}$, and taking into account the mobile sink, we obtain by comparison $ba^m = ba^{t-1}$ and so $m = t - 1$, a contradiction, since a^t is a suffix of w. Therefore $v' = a^t$ and so is independent from z as required. $\square$

Corollary 4.11 *Let $H \leq_{f.g.} F_2$ and $t > \frac{1}{2}\zeta(H)$. Then the set*

$$\mathcal{X}(t, H) = \{\mathcal{A}_t(K) \mid K \in \Sigma^*(H)\}$$

is finite and effectively constructible.

Proof. By Lemma 4.7, every automaton $\mathcal{A}(K)$, $K \in \Sigma^*(H)$, has at most $\sigma(H)$ singularities. By definition of a t-truncation, every state in $\mathcal{A}_t(K)$

is at distance at most t from a singularity, and hence the size of $\mathcal{A}_t(K)$ is bounded. Thus $\mathcal{X}(t, H)$ is finite.

The proof of Theorem 4.10 provides a straighforward algorithm to compute all its elements. Indeed, all we need is to compute the finite sets

$$\mathcal{X}_n(t, H) = \{\mathcal{A}_t(K) \mid K \in \Sigma^n(H)\}$$

until reaching

$$\mathcal{X}_{n+1}(t, H) \subseteq \bigcup_{i=0}^{n} \mathcal{X}_i(t, H), \tag{3}$$

which must occur eventually since $\mathcal{X}(t, H) = \cup_{i \geq 0} \mathcal{X}_i(t, H)$ is finite. Why does (3) imply $\mathcal{X}(t, H) = \cup_{i \geq 0}^{n} \mathcal{X}_i(t, H)$? Suppose that $\mathcal{B} \in \mathcal{X}_m(t, H) \setminus (\cup_{i \geq 0}^{n} \mathcal{X}_i(t, H))$ with m minimal, say $\mathcal{B} = \mathcal{A}_t(\varphi(K))$ with $K \in \Sigma^{m-1}(H)$ and $\varphi \in \Sigma$. By minimality of m, we have $\mathcal{A}_t(K) \in \cup_{i \geq 0}^{n} \mathcal{X}_i(t, H)$. Thus $\mathcal{A}_t(K) = \mathcal{A}_t(K')$ for some $K' \in \cup_{i=0}^{n} \Sigma^i(H)$. Now Theorem 4.10 yields

$$\mathcal{B} = \mathcal{A}_t(\varphi(K)) = \mathcal{A}_t(\varphi(K')) \in \bigcup_{i=0}^{n+1} \mathcal{X}_i(t, H) = \bigcup_{i=0}^{n} \mathcal{X}_i(t, H),$$

a contradiction. Therefore $\mathcal{X}(t, H) = \cup_{i \geq 0}^{n} \mathcal{X}_i(H)$ as claimed. $\square$

5 Back to orbit problems in F_2

We saw in Section 2 that it is decidable whether a given element $u \in F_2$ has an automorphic image in a given rational subset of F_2, and in particular in a given finitely generated subgroup of F_2 (Corollary 2.3 above). We use truncated automata to give a different proof of this result in the finitely generated subgroup case. We also prove the decidability of mixed orbit problems under the action of certain rational subsets of Σ^*.

5.1 Some mixed orbit problems

The archetypal result in this section is the solution of the following mixed orbit problem.

Proposition 5.1 *Let $u \in F_2$ an $H \leq_{fg} F_2$. The set of automorphisms $\varphi \in \Sigma^*$ such that $u \in \varphi(H)$ (resp. a conjugate of u lies in $\varphi(H)$), is an effectively constructible rational subset of Σ^*.*

Proof. Let $t > \max(\frac{1}{2}\zeta(H), \frac{1}{2}|u|)$: if $\varphi \in \Sigma^*$, then $u \in \varphi(H)$ if and only if u labels a loop at the origin in $\mathcal{A}(\varphi(H))$. Note that this is the case if and only if u labels a loop at the origin in $\mathcal{A}_t(\varphi(H))$.

We now view Σ as a finite alphabet (besides being a subset of $\mathsf{Aut}\,F_2$) and we consider the Σ-transition system $\mathcal{B}_t(H)$ defined as follows. (A Σ-*transition system* is defined like a Σ-automaton, omitting the specification of the initial and terminal states.) The state set of $\mathcal{B}_t(H)$ is $\mathcal{X}(t, H)$ (see Corollary 4.11) and its transitions are the triples $\mathcal{A}_t(K)\overset{\varphi}{\longrightarrow}\mathcal{A}_t(\varphi(K))$, for each $\mathcal{A}_t(K) \in \mathcal{X}(t, H)$ and $\varphi \in \Sigma$. Note that $\mathcal{X}(t, H)$ is finite and effectively constructible by Corollary 4.11 and the transitions of $\mathcal{B}_t(H)$ are well-defined by Theorem 4.10. Moreover, this transition system is complete and deterministic by construction (and so defines recognizable subsets of Σ^* as a submonoid of $\mathsf{Aut}\,F_2$). It is immediate that if the word $(\varphi_1, \ldots, \varphi_n) \in \Sigma^*$ labels a path from $\mathcal{A}_t(K)$ to $\mathcal{A}_t(K')$ in $\mathcal{B}_t(H)$ $(K, K' \in \Sigma^*(H))$, then $\mathcal{A}_t(K') = \mathcal{A}_t(\varphi_n \ldots \varphi_1(K))$.

Now consider the automaton formed by the transition system $\mathcal{B}_t(H)$ with initial state $\mathcal{A}_t(H)$ and terminal states the elements $\mathcal{A}_t(K)$ of $\mathcal{X}(t, H)$ such that u labels a loop at the origin in $\mathcal{A}_t(K)$ (i.e. $u \in K$). The above discussion shows that the language accepted by this automaton is the set of words $(\varphi_1, \ldots, \varphi_n) \in \Sigma^*$ such that $u \in \varphi_n \cdots \varphi_1(H)$. Thus the set of all $\varphi \in \Sigma^*$ such that $u \in \varphi(H)$ is rational and effectively constructible.

Observe that a conjugate of u lies in $\varphi(H)$ $(\varphi \in \mathsf{Aut}\,F_2)$, if and only if u labels a loop at the origin in $\mathcal{A}(\lambda_w \varphi(H))$ for some $w \in F_2$, if and only if $cc(u)$ labels a loop somewhere in $\mathcal{A}(\varphi(H))$. We now consider the Σ-transition system $\mathcal{B}_t(H)$ as above, with the same initial state, and we take as terminal states the elements $\mathcal{A}_t(K)$ of $\mathcal{X}(t, H)$ such that $cc(u)$ labels a loop anywhere in $\mathcal{A}_t(K)$. The language in Σ^* accepted by the resulting automaton is the set of $\varphi \in \Sigma^*$ such that $\varphi(H)$ contains a conjugate of u. $\square$

The same idea — and the same transition system— can be used to algorithmically solve a number of other orbit problems.

Theorem 5.2 *Let* $H, K \leq_{fg} F_2$, $u, u_1, \ldots, u_k \in F_2$ *and* $R \in \mathsf{Rat}\,\Sigma^*$. *Then the following problems are decidable:*

(1) *whether* $u \in \mu(H)$ *for some* $\mu \in R$;

(1') *whether a conjugate of* u *lies in* $\mu(H)$ *for some* $\mu \in R$; *that is, whether* $u \in \mu(H)$ *for some* $\mu \in \Lambda R$;

(2) *whether* $K \subseteq \mu(H)$ *for some* $\mu \in R$;

(2′) *whether a conjugate of K is contained in $\mu(H)$ for some $\mu \in R$; that is, whether $K \subseteq \mu(H)$ for some $\mu \in \Lambda R$;*

(3) *whether $K = \mu(H)$ for some $\mu \in R$;*

(3′) *whether a conjugate of K is equal to $\mu(H)$ for some $\mu \in R$; that is, whether $K = \mu(H)$ for some $\mu \in \Lambda R$;*

(4) *whether there exist $w_1, \ldots, w_k \in F_2$ such that $\lambda_{w_1}(u_1), \ldots, \lambda_{w_k}(u_k) \in \mu(H)$ for some $\mu \in R$.*

In addition, for each of these problems, the set of morphisms $\mu \in \Sigma^*$ that it defines is rational and effectively constructible.

Proof. The solutions of Problems (1) and (1′) follow from Proposition 5.1: the set X of automorphisms $\varphi \in \Sigma^*$ such that $\varphi(H)$ contains u (resp. a conjugate of u) is rational and we can compute a Σ-automaton recognizing that set. Since $\mathcal{B}_t(H)$ is deterministic and complete, we only have to decide whether X has a non-empty intersection with the given rational set R, a classical decidable result from automata theory.

The other proofs follow the same pattern, and correspond to variants of Proposition 5.1. Let us consider Problem (2) and let $u_1, \ldots, u_k$ be generators of K. Then we need to consider the Σ-transition system $\mathcal{B}_t(H)$ with $t > \max(\frac{1}{2}\zeta(H), \frac{1}{2}|u_1|, \ldots, \frac{1}{2}|u_k|)$, and to choose as terminal states the elements $\mathcal{A} \in \mathcal{X}(t, H)$ such that $\bar{u}_1, \ldots, \bar{u}_k$ label loops at the origin in $\mathcal{A}$.

For Problem (2′), we consider a *cyclically reduced* conjugate K' of K, that is, one such that the origin in $\mathcal{A}(K')$ has degree at least 2 (if the origin in $\mathcal{A}(K)$ has degree 1, choose any vertex v with degree at least 2 as the new origin and let K' be the corresponding conjugate). Let $u_1, \ldots, u_k$ be generators of K'. Then a conjugate of K lies in $\mu(H)$ if and only if the $\bar{u}_i$ label loops around the same vertex of $\mathcal{A}(\mu(H))$. Thus it suffices to choose $t > \max(\frac{1}{2}\zeta(H), \frac{1}{2}|u_1|, \ldots, \frac{1}{2}|u_k|)$, and to take as terminal states the elements $\mathcal{A} \in \mathcal{X}(t, H)$ such that $\bar{u}_1, \ldots, \bar{u}_k$ label loops around the same vertex of $\mathcal{A}$.

For Problem (3), we choose again $t > \max(\frac{1}{2}\zeta(H), \frac{1}{2}|u_1|, \ldots, \frac{1}{2}|u_k|)$, where $u_1, \ldots, u_k$ are generators of K. In particular, t is large enough to have $\mathcal{A}_t(K) = \mathcal{A}(K)$, and we choose a single terminal state, $\mathcal{A}_t(K)$ (if $\mathcal{A}_t(K) \in \mathcal{X}(t, H)$; if that is not the case, then Problem (3) is decidable, in the negative). Then we have an automaton which recognizes the set $L(t, K)$ of all $\mu \in \Sigma^*$ such that $\mathcal{A}_t(\mu(H)) = \mathcal{A}_t(K) = \mathcal{A}(K)$. Observe now that truncation creates (pairs of) degree 1 vertices: the automorphisms $\mu \in L(t, K)$ are such that $\mathcal{A}_t(\mu(H))$ has at most one degree 1 vertex (the origin), and

hence $\mathcal{A}_t(\mu(H)) = \mathcal{A}(\mu(H))$. Thus our automaton recognizes the set of all $\mu \in \Sigma^*$ such that $\mathcal{A}(\mu(H)) = \mathcal{A}(K)$, that is, such that $\mu(H) = K$.

For Problem $(3')$, we consider a cyclically reduced conjugate K' of K and an integer t as in Problem $(2')$. Again, we have $\mathcal{A}_t(K') = \mathcal{A}(K')$. We choose as terminal states the elements of $\mathcal{X}(t, H)$ of the form $\mathcal{A}_t(\lambda_w(K))$ $(w \in F_2)$. These automata are of one of the following types:

| form (a) | form (b) | form (c) |
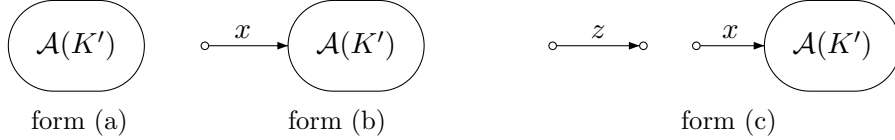
with $|x| \le t$ and $|z| = t$. As in the discussion of Problem (3), the existence of a μ-labeled path in $\mathcal{B}_t(H)$ from $\mathcal{A}_t(H)$ to an automaton of type (a) or (b) shows that $\mu(H)$ is a conjugate of K', and hence of K. If the path in $\mathcal{B}_t(H)$ ends in an automaton of type (c), then $\mu(H)$ is a conjugate of K' of the form $zyxwK'(zyxw)^{-1}$ or $z^{-1}yxwK'(z^{-1}yxw)^{-1}$ for some y, w such that $zyxw$ or $z^{-1}yxw$ is reduced. We then conclude the proof of the decidability of Problem $(3')$ as usual.

Finally, for Problem (4), we choose $t > \max(\frac{1}{2}\zeta(H), \frac{1}{2}|cc(u_1)|, \dots, \frac{1}{2}|cc(u_k)|)$ and we choose as terminal states the elements $\mathcal{A} \in \mathcal{X}(t, H)$ such that each $cc(u_i)$ $(i = 1, \dots, k)$ labels a loop at some vertex in $\mathcal{A}$. $\square$

We can also consider finitely many subgroups H_i in (4) and many other variations.

A simple rewriting of Theorem 5.2 in terms of orbit problems (see the introduction) yields the following corollary.

Corollary 5.3 *Let $H \le_{fg} F_2$, $u \in F_2$ and $R \in \mathsf{Rat}\,\Sigma^*$. Then it is decidable whether the orbit of u under the action of R^{-1} (resp. ΛR^{-1}) meets H.*

If in addition $K \le_{fg} F_2$, then it is decidable whether H contains an element of the orbit of K under the action of R^{-1} or ΛR^{-1}; whether H is contained in an element of the orbit of K under the action of R or ΛR; and whether K is an element of the orbit of H under the action of R, R^{-1}, ΛR or ΛR^{-1}.

Applying Corollary 5.3 to the case where u is a letter in A, we get a statement about primitive elements.

Corollary 5.4 *Let $H \le_{fg} F_2$ and $R \in \mathsf{Rat}\,\Sigma^*$. Then it is decidable whether H contains a primitive element of the form $\mu(a)$, $\mu^{-1} \in R$ (resp. $\mu^{-1} \in \Lambda R$).*

Remark 5.5 Let S be a subset of R_2 such that, for each rational set S', one can decide whether $S \cap S'$ is empty or not. Then Problems (1'), (2') and (3') in Theorem 5.2 are decidable even if we restrict the conjugating factors to be in S, that is, if we replace Λ by $\{\lambda_s \mid s \in S\}$ in the statement of these problems. The same restriction can be imposed to Λ in the statements of Corollaries 5.3 and 5.4.

Similarly, Problem (4) in Theorem 5.2 remains decidable even if we require the w_i to be in fixed subsets S_i ($i = 1, \ldots, k$) such that, for each rational set S', one can decide whether $S_i \cap S'$ is empty or not. $\square$

5.2 Orbits under invertible substitutions

Invertible substitutions are an interesting special case of the rational subsets of $\mathsf{Aut}\, F_2$ discussed in Section 5.1. This leads to the following statement.

Corollary 5.6 *The problems discussed in Theorem 5.2 and Corollaries 5.3 and 5.4 are decidable also if R is assumed to be a rational subset of $IS(F_2)$ or $IS(F_2)^{-1}$.*

Proof. If $R \in \mathsf{Rat}\, IS(F_2)$, then $R \in \mathsf{Rat}\, \Sigma^*$ by Lemma 3.5, and we simply apply Theorem 5.2 and Corollaries 5.3 and 5.4.

If $R \in IS^{-1}(F_2)$, then $R = \varphi_{a,b^{-1}}\, R'\, \varphi_{a,b^{-1}}$ for some $R' \in \mathsf{Rat}\, \Sigma^*$ by Lemma 3.5 (R' is the set of inverses of the elements of R). Problem (1) in Theorem 5.2 on instance u, H and R, for example, is equivalent to the same problem on instances $\varphi_{a,b^{-1}}(u)$, $\varphi_{a,b^{-1}}(H)$ and R', which we know to be decidable. The other problems are handled in the same fashion. $\square$

5.3 Another solution of the mixed orbit problem for $\mathsf{Aut}\, F_2$

Our proof relies on truncated automata and Theorem 3.4. The key is to bound the powers of $\varphi_{a,ba}$ that we need to consider, and is achieved in view of our previous bound for the length of homogeneous cycles.

Let $u \in F_2$ and $H \leq_{f.g.} F_2$. We want to show that it is decidable whether $\mu(u) \in H$ for some $\mu \in \mathsf{Aut}\, F_2$. By Theorem 3.4, and since $\Psi^{-1} = \Psi$, it suffices to decide whether there exist $w \in F_2$ and $n \geq 0$ such that one of the following conditions hold:

- $\lambda_w \varphi_{a,ba}^n \varphi_{a^{-1},b}(u) \in \Sigma_0^* \Psi(H)$;

- $\lambda_w \varphi_{a,ba}^n \varphi_{a^{-1},b^{-1}}(u) \in \Sigma_0^* \Psi(H)$.

Since Ψ is finite, it suffices to be able to decide whether

there exist $w \in F_2$, $n \geq 0$ and $\mu \in \Sigma_0^*$ such that $\lambda_w \varphi_{a,ba}^n(u) \in \mu(H)$. (4)

We start by considering the case $n = 0$. By Proposition 3.1 (i), we may replace $\lambda_w \varphi_{a,ba}^n$ by $\varphi_{a,ba}^n \lambda_w$, so we may assume that u is cyclically reduced. And by Proposition 1.2 (iv), our problem further reduces to asking if one can decide whether

u labels a loop in $\mathcal{A}(\mu(H))$ for some $\mu \in \Sigma_0^*$. (5)

We note that every loop contains either the origin or a singularity: if it does not contain the origin, then there is a path from the origin to a state in the loop, and the first contact between that path and the loop is a source or a sink. Now let us fix $t > \max(\frac{1}{2}\zeta(H), \frac{1}{2}|u|)$: then u labels a loop in $\mathcal{A}(\mu(H))$ if and only if u labels a loop in $\mathcal{A}_t(\mu(H))$. By the appropriate variant of Corollary 4.11 (where Σ is replaced with Σ_0^*) we can effectively compute the finite set

$$\mathcal{X}_0(H) = \{\mathcal{A}_t(K) \mid K \in \Sigma_0^*(H)\}.$$

Thus (5) is decidable, and hence (4) is decidable for $n = 0$. It is also decidable for any fixed n (applying the case $n = 0$ to $\varphi_{a,ba}^n(u)$ instead of u).

We now consider (4) in its full generality. If $u \in R_a$, then we are reduced to the case $n = 0$ since $\varphi_{a,ba}(u) = u$. So we assume that b or b^{-1} occurs in u, and by conjugation again, we may assume that u starts with b or ends with b^{-1} (and not both since u is cyclically reduced).

Let M be the least common multiple of $1, 2, \ldots, \delta_0(H)$. In order to prove (4), it suffices to show that

if there exist $w \in F_2$, $n \geq 0$ and $\mu \in \Sigma_0^*$ such that $\lambda_w \varphi_{a,ba}^n(u) \in \mu(H)$, then there exists such a triple (w, n, μ) with $n < |u| + \max(M, \delta(H))$.

Since we have proved (4) for bounded n, the latter property is decidable, and hence (4) is decidable in general.

So we are left with the task of proving this reduced claim. Let (w, n, μ) be such that $\lambda_w \varphi_{a,ba}^n(u) \in \mu(H)$, with n minimal, and let us suppose that $n \geq |u| + \max(M, \delta(H))$.

Write $u = a^{i_0} b^{\varepsilon_1} a^{i_1} \ldots b^{\varepsilon_k} a^{i_k}$ with $k \geq 1$ and $\varepsilon_\ell = \pm 1$ for every ℓ. If $m \geq 0$, then

$$\varphi_{a,ba}^m(u) = \varphi_{a,ba^m}(u) = a^{j_0} b^{\varepsilon_1} a^{j_1} \ldots b^{\varepsilon_k} a^{j_k}$$

with

$$j_\ell = \begin{cases} i_\ell + m & \text{if } \varepsilon_\ell = \varepsilon_{\ell+1} = 1, \text{ or } \ell = k \text{ and } \varepsilon_k = 1 \\ i_\ell - m & \text{if } \varepsilon_\ell = \varepsilon_{\ell+1} = -1, \text{ or } \ell = 0 \text{ and } \varepsilon_1 = -1 \\ i_\ell & \text{in all other cases.} \end{cases}$$

Recall that u is cyclically reduced, and that it starts with b ($i_0 = 0$ and $\varepsilon_1 = 1$) or ends with b^{-1} ($i_k = 0$ and $\varepsilon_k = -1$). It follows that $\varphi_{a,ba^m}(u)$ is cyclically reduced and that it too starts with b or ends with b^{-1}.

By Proposition 1.2 (iv), $\varphi_{a,ba}^n(u)$ labels a loop α in $\mathcal{A}(\mu(H))$. Moreover, we have

$$\varphi_{a,ba}^n(u) = a^{r_0} b^{\varepsilon_1} a^{r_1} \dots b^{\varepsilon_k} a^{r_k}, \quad \varphi_{a,ba}^{n-M}(u) = a^{s_0} b^{\varepsilon_1} a^{s_1} \dots b^{\varepsilon_k} a^{s_k},$$

with

$$\begin{cases} r_\ell = i_\ell + n, \ s_\ell = r_\ell - M & \text{if } \varepsilon_\ell = \varepsilon_{\ell+1} = 1, \text{ or } \ell = k \text{ and } \varepsilon_k = 1 \\ r_\ell = i_\ell - n, \ s_\ell = r_\ell + M & \text{if } \varepsilon_\ell = \varepsilon_{\ell+1} = -1, \text{ or } \ell = 0 \text{ and } \varepsilon_1 = -1 \\ s_\ell = r_\ell = i_\ell & \text{in all other cases.} \end{cases}$$

In the first and second cases, $|r_\ell| > n - |u| \geq \max(M, \delta(H))$; and in the last case, $|r_\ell| < |u|$. Thus, for the indices ℓ such that $r_\ell \neq s_\ell$, we have $r_\ell > \delta(H)$. We now show that the fragments of the loop α labeled by the factors a^{r_ℓ} such that $r_\ell \neq s_\ell$, fail to be cycle-free in $\mathcal{A}(\mu(H))$.

Recall that $\mu \in \Sigma_0^*$. If $\mu = \text{id}$ or $\varphi_{b^{-1},a^{-1}}$, the result is immediate since $r_\ell > \delta(H) = \delta(\mu(H))$. If $\mu = \varphi_{a,ba}\nu$ with $\nu \in \Sigma_0^*$, then we can use Proposition 3.1 (i) to reduce n, a contradiction. Hence we may assume that $\mu = \varphi_{b^{-1},a^{-1}}\nu$ with $\nu \in \Sigma_0^*$, $\nu \neq \text{id}$. Since $\varphi_{b^{-1},a^{-1}}^2 = \text{id}$, we may further assume that $\mu = \varphi_{b^{-1},a^{-1}}\varphi_{a,ba}\nu'$ with $\nu' \in \Sigma_0^*$. Then the vertices involved in the a^{r_ℓ}-labeled fragment of α form a path in $\mathcal{A}(\varphi_{a,ba}\nu'(H))$ labeled b^{r_ℓ}. Since $r_\ell > \delta(H)$, we also have $r_\ell > \sigma(H) \geq \sigma(\nu'(H))$ (Lemma 4.7), and hence this path is not cycle-free by Lemma 4.6.

So, for each ℓ such that $r_\ell \neq s_\ell$, the fragment of α labeled by the factor a^{r_ℓ} of $\varphi_{a,ba}^n(u)$ fails to be cycle-free, and must be read along a cycle of $\mathcal{A}(\mu(H))$ (in an inverse automaton, if a homogeneous path contains a cycle, then it reads entirely along that cycle).

By definition, M is a multiple of the length c_ℓ of that cycle. Now compare $\varphi_{a,ba}^{n-M}(u)$ and $\varphi_{a,ba}^n(u)$: wherever the a-factors a^{r_ℓ} and a^{s_ℓ} are different, their difference is either a^M or a^{-M}, and hence it consists of a whole number of passages around the length c_ℓ cycle. Therefore $\varphi_{a,ba}^{n-M}(u)$ labels a path in $\mathcal{A}(\mu(H))$ as well. This contradicts the minimality of n and completes the proof.

Remark 5.7 The *a priori* complexity of the algorithms discussed in Section 5 is very high: if u has length at most n and $\mathcal{A}(H)$ has at most n states, then $\sigma(H), \zeta(H) \leq n$ and the truncated automata can have exponentially many states. There can therefore be super-exponentially many truncated automata, forming the states of the transition system $\mathcal{B}_t(H)$ – in which we must solve a reachability problem (polynomial in the number of states of the transition system).

References

[1] M. Benois, Parties rationnelles du groupe libre, *C. R. Acad. Sci. Paris* 269 (1969), 1188–1190.

[2] J. Berstel, A. Lauve, C. Reutenauer and F. Saliola. *Combinatorics on words: Christoffel words and repetitions in words*, CRM monograph series 27, AMS, 2009.

[3] P. Brinkmann, Detecting automorphic orbits in free groups, `arXiv:0806.2889v1`.

[4] A. Clifford and R. Goldstein, Subgroups of free groups and primitive elements, *J. Group Theory*, to appear.

[5] M. Cohen, W. Metzler and A. Zimmermann. What does a basis of $F(a, b)$ look like?, *Math. Ann.* 257 (1981), 435–445.

[6] V. Diekert, C. Gutiérrez and C. Hagenah, The existential theory of equations with rational constraints in free groups is PSPACE-complete, *Information and Computation* 202 (2005), 105–140.

[7] S. Gersten, On Whitehead's algorithm, *Bull. Am. Math. Soc.* 10 (1984) 281–284.

[8] , R. Goldstein, An algorithm for potentially positive words in F_2, In *Combinatorial group theory, discrete groups, and number theory*, volume 421 of *Contemp. Math.*, pages 157-168. Amer. Math. Soc., 2006.

[9] J. E. Hopcroft and J. D. Ullman, *Introduction to Automata Theory, Languages and Computation*, Addison-Wesley, 1979.

[10] I. Kapovich and A. Myasnikov, Stallings foldings and subgroups of free groups, *J. Algebra* 248 (2002), 608–668.

[11] D. Lee, On several problems about automorphisms of the free group of rank two, *J. Algebra* 321(1) (2009), 167-193.

[12] M. Lothaire. *Algebraic combinatorics on words*, Encyclopedia of Mathematics and its Applications, vol. 90, Cambridge University Press, 2002.

[13] R. C. Lyndon and P. E. Schupp, *Combinatorial Group Theory*, Springer-Verlag 1977.

[14] G. S. Makanin, Equations in a free group (Russian), *Izv. Akad. Nauk. SSSR* Ser. Mat. 46 (1983), 1199–1273; English translation in *Math. USSR Izv.* 21 (1983).

[15] A. Miasnikov, E. Ventura and P. Weil, Algebraic extensions in free groups, in *Algebra and Geometry in Geneva and Barcelona* (G.N. Arzhantseva, L. Bartholdi, J. Burillo and E. Ventura eds.), Trends in Mathematics, Birkhaüser (2007), pp. 225–253.

[16] J. Nielsen. Die Isomorphismen der allgemeinen unendlichen Gruppe mit zwei Erzeugenden, *Math. Ann.* 78 (1918), 385–397.

[17] R. P. Osborne and H. Zieschang. Primitives in the free group on two generators, *Invent. Math.* 63 (1981), 17–24.

[18] A. Roig, E. Ventura and P. Weil, On the complexity of the Whitehead minimization problem, *Int. J. Alg. Comput.* 17 (2007), 1611-1634.

[19] J. Rotman. *An introduction to the theory of groups*, 4th edition, Springer, 1995.

[20] J.-P. Serre. *Arbres, amalgames, SL_2*, Astérisque 46, Soc. Math. France, 1977. English translation: *Trees*, Springer Monographs in Mathematics, Springer, 2003.

[21] A. Shenitzer, Decomposition of a group with a single defining relation into a free product, *Proc. Amer. Math. Soc.* 6 (1955), 273-279.

[22] V. Shpilrain. Recognizing automorphisms of the free groups, *Arch. Math.* 62 (1994), 385-392.

[23] P. V. Silva and P. Weil. On an algorithm to decide whether a free group is a free factor of another, *RAIRO Theoretical Informatics and Applications* 42 (2008), 395–414.

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

[25] Z. X. Wen and Z. Y. Wen. Local isomorphisms of invertible substitutions, *C. R. Acad. Sci. Paris Sér. I Math.* 318 (1994), 299–304.

[26] J.H.C. Whitehead, On equivalent sets of elements in a free group, *Annals of Mathematics* 37 (1936) 782–800.