



HAL
open science

On an algorithm to decide whether a free group is a free factor of another

Pedro Silva, Pascal Weil

► **To cite this version:**

Pedro Silva, Pascal Weil. On an algorithm to decide whether a free group is a free factor of another. *RAIRO - Theoretical Informatics and Applications (RAIRO: ITA)*, 2008, 42, pp.395-414. hal-00096810v2

HAL Id: hal-00096810

<https://hal.science/hal-00096810v2>

Submitted on 16 Nov 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ON AN ALGORITHM TO DECIDE WHETHER A FREE GROUP IS A FREE FACTOR OF ANOTHER *

PEDRO V. SILVA¹ AND PASCAL WEIL²

Abstract. We revisit the problem of deciding whether a finitely generated subgroup H is a free factor of a given free group F . Known algorithms solve this problem in time polynomial in the sum of the lengths of the generators of H and exponential in the rank of F . We show that the latter dependency can be made exponential in the rank difference $\text{rank}(F) - \text{rank}(H)$, which often makes a significant change.

Résumé. Nous revenons sur la question de décider si un sous-groupe finiment engendré H est facteur libre d'un groupe libre donné F . On trouve dans la littérature des algorithmes qui résolvent ce problème en temps polynomial en la somme des longueurs des générateurs de H , et exponentiel en le rang de F . Nous montrons que l'on peut remplacer la dépendance exponentielle en $\text{rank}(F)$ par une dépendance exponentielle en la différence $\text{rank}(F) - \text{rank}(H)$, ce qui change souvent les choses de façon considérable.

1991 Mathematics Subject Classification. 20E05,05C25.

The combinatorial aspects of group theory have attracted the attention of theoretical computer scientists for a long time, and for a variety of reasons. There is no need to recall the importance of the concept of monoid (e.g. free, finite) in the theory of automata since the foundational results of Schützenberger and Eilenberg (see the books [1, 18]), and groups form a special case of monoids that

Keywords and phrases: combinatorial group theory, free groups, free factors, inverse automata, algorithms

* *The first author acknowledges support from C.M.U.P., financed by F.C.T. (Portugal) through the programmes POCTI and POSI, with national and European Community structural funds. Both authors acknowledge support from the European Science Foundation program AutoMathA.*

¹ Centro de Matemática, Faculdade de Ciências – Universidade do Porto – R. Campo Alegre 687 – 4169-007 Porto, Portugal; email: pvsilva@fc.up.pt

² LaBRI, CNRS – 351 cours de la Libération – 33405 Talence Cedex – France; email: pascal.weil@labri.fr

© EDP Sciences 1999

sometimes play an important role in purely language- and monoid-theoretic questions (e.g. the type II conjecture, see [6] for a survey). Algorithmic questions (the word problem, the conjugacy problem, *etc.*) have been very influential in group theory throughout the 20th century, starting from the work of Dehn, and specialists of combinatorics on words find a particular interest in the challenges posed by the analogous combinatorics of the free group. Recent work on important non-commutative groups like the Thompson group and the so-called automata groups strongly relies on the formalism of finite state automata (see [3, 7, 22] for recent examples). It is already a classical result that these automata-theoretic and combinatorial points of view converge in the (admittedly simpler) study of the subgroups of free groups, this is central in the algorithmic problem tackled here and is discussed in detail in the first part of this paper.

Let us also mention another reason for the recent multiplication of research projects on the boundary between computer science and combinatorial group theory. Public-key cryptography relies heavily on group theory: sometimes finite groups such as the groups of units in modular arithmetic, or the groups of rational points on elliptic curves over finite fields, sometimes infinite non-commutative groups like the braid groups (see for instance [2, 4, 21] and many others). At any rate, the design of more robust cryptographic schemes and the attack of such schemes rely on a deeper understanding of the combinatorial and algorithmic properties of non-commutative groups.

As mentioned above, the combinatorial and algorithmic problems concerning free groups are of special interest. Free groups are archetypal groups, whose structure is far from being totally elucidated, and the efficient solution of standard problems in their context can shed some light on the possible solution of the same problems in more complex groups. Moreover, the solution of algorithmic problems in free groups may be more attainable since we can use the resources of combinatorics on words and automata theory.

We now present the specific algorithmic problem addressed in this paper. For the classical facts about free groups recorded below without a reference, we refer the reader to the book by Lyndon and Schupp [12]. It is well-known that the minimal sets of generators, or *bases*, of a free group F all have the same cardinality, called the *rank* of F . Moreover, if F has finite rank r , every r -element generating set of F is a basis, see [12, Prop. I.3.5]. In this paper, we consider only finite rank free groups.

Let H be a subgroup of a free group F , written $H \leq F$. Then H itself is a free group whose rank may be greater than the rank of F . We say that H is a *free factor* of F , written $H \leq_{\text{ff}} F$, if there exist bases B of H and A of F such that $B \subseteq A$ (free factors can be defined in all groups by a universal property, but the operational definition given here is sufficient for the purpose of this study). It is well known that one can decide whether a given finite rank subgroup $H \leq F$ is a free factor of F , but the known algorithms have a rather high time complexity. More precisely, the best of these algorithms require time that is polynomial in the size of H and exponential in the rank of F , see Section 1.3 below for the details.

Here, the size of H is taken to be the sum of the lengths of a finite set of generators of H in F .

We propose a new algorithm to decide whether a given finitely generated subgroup H is a free factor of the free group F , which is polynomial in the size of H and exponential in the rank difference between F and H . In many instances, this represents a substantial advantage over exponential dependency in the rank of F .

Our algorithm relies essentially on a careful analysis of the construction of the *graph representation* of H . More precisely, once a basis A of the ambient free group F is fixed, there is a natural and elegant representation of the finitely generated subgroups of F by A -labeled graphs (or inverse automata). This construction — a graphical representation of ideas that go back to the early part of the twentieth century [19, Chap. 11] — was made explicit by Serre [20] and Stallings [23]. It has been used to great profit by many authors since the late 1970s, see [8, 13, 14, 26] for recent examples. Given a finite set of generators of H (as reduced words over the alphabet $A \cup A^{-1}$), the graph representation of H can be effectively constructed (see [23], [14], *etc*). The number of vertices and edges of this graph is bounded above by ℓ , the sum of the lengths of a set of generators of H , and the whole representation can be computed in time at most $O(\ell^2)$ (in fact, in time $O(\ell \log^* \ell)$ according to a recent announcement¹ by Touikan [25]). We discuss this representation in more detail in Section 1.2 below, and we show in Sections 2 and 3 how to use it to decide more efficiently the free factor relation.

It is interesting to note that our algorithm is the first to be expressed entirely in terms of the graph representation of H . Let us also emphasize that we do not claim that our algorithm is optimal. It is an open question whether one can decide the free factor relation $H \leq_{\text{ff}} K$ in time polynomial both in the size of H and in the size, or the rank of K .

1. BACKGROUND

If A is a basis of a free group F , we often write $F = F(A)$ and we represent the elements of F as reduced words over the alphabet A . More precisely, we consider the set of all words on the symmetrized alphabet $\tilde{A} = A \cup A^{-1}$, where $A^{-1} = \{a^{-1} \mid a \in A\}$ is a set that is disjoint from A , equipped with an explicit bijection with A , namely $a \mapsto a^{-1}$. It is customary to extend the mapping $u \mapsto u^{-1}$ to all words $u \in \tilde{A}^*$ by letting $(a^{-1})^{-1} = a$ for each $a \in A$, $1^{-1} = 1$ (where 1 denotes the empty word) and $(a_1 a_2 \cdots a_n)^{-1} = a_n^{-1} \cdots a_2^{-1} a_1^{-1}$ for all $a_1, \dots, a_n \in \tilde{A}$. A word in \tilde{A}^* is *reduced* if it contains no factor of the form aa^{-1} or $a^{-1}a$ with $a \in A$, and it is well known that F can be identified with the set of reduced words over A . We denote by ρ the map that assigns to each word u the corresponding reduced word $u\rho \in F(A)$, obtained by iteratively deleting all factors of the form aa^{-1} or $a^{-1}a$ ($a \in A$).

¹For a positive integer n , $\log^*(n)$ is the least integer k such that the k -th iterate of the log function of n is at most 1. The growth of $\log^*(n)$ is so slow that it can be considered a constant for all practical purposes...

1.1. ON INVERSE AUTOMATA

We describe the main tool for the representation of subgroups of free groups in terms of automata (see [17]). Readers less familiar with this terminology may think of automata as edge-labeled directed graphs.

An *automaton* on alphabet A is a triple of the form $\mathcal{A} = (Q, q_0, E)$ where Q is a finite set called the *state set*, $q_0 \in Q$ is the *initial state*, and $E \subseteq Q \times A \times Q$ is the set of *edges*, or *transitions*. A transition (p, a, q) is said to be from state p , to state q , with label a . The *label of a path* in \mathcal{A} (a finite sequence of consecutive transitions) is the sequence of the labels of its transitions, a word on alphabet A , that is, an element of the free monoid A^* . We write $p \xrightarrow{u} q$ if there is a path from state p to state q with label u . The language accepted by \mathcal{A} is the set $L(\mathcal{A})$ of all words in A^* which label a path in \mathcal{A} from q_0 to q_0 .

This definition of automata leads naturally to the definition of a *homomorphism* φ from an automaton $\mathcal{A} = (Q, q_0, E)$ to an automaton $\mathcal{A}' = (Q', q'_0, E')$ (over the same alphabet A): φ is a mapping from Q to Q' such that $\varphi(q_0) = q'_0$, and such that whenever $(p, a, q) \in E$, we also have $(\varphi(p), a, \varphi(q)) \in E'$. The homomorphism φ is an *isomorphism* if it is a bijection and if φ^{-1} is also a homomorphism.

The automaton \mathcal{A} is called *deterministic* if no two distinct edges with the same initial state bear the same label, that is,

$$(p, a, q), (p, a, q') \in E \implies q = q'.$$

The automaton is called *trim* if every state $q \in Q$ lies in some path from q_0 to q_0 .

In the sequel, we consider automata where the alphabet is symmetrized, that is, the alphabet is of the form $\tilde{A} = A \cup A^{-1}$. We say that \mathcal{A} is *dual* if for each $a \in A$, there is an a -labeled edge from state p to state q if and only if there is an a^{-1} -labeled edge from q to p ,

$$(p, a, q) \in E \iff (q, a^{-1}, p) \in E.$$

Let us immediately record the following fact.

Fact 1.1. Let \mathcal{A} be a deterministic dual automaton. If a word u labels a path in \mathcal{A} from state p to state q , then so does the corresponding reduced word $u\rho$. Moreover $L(\mathcal{A})$ is a submonoid of \tilde{A}^* and $L(\mathcal{A})\rho$ is a subgroup of $F(A)$.

Now let $\mathcal{A} = (Q, q_0, E)$ be a trim dual automaton and let $p, q \in Q$ be states of \mathcal{A} . If $w = a_1 \cdots a_n \in \tilde{A}^*$ is a non-empty word, the expansion of \mathcal{A} by (p, w, q) is the automaton obtained from \mathcal{A} by adding $n - 1$ vertices q_1, \dots, q_{n-1} and $2n$ edges

$$p \xrightarrow{a_1} q_1 \xrightarrow{a_2} \dots \xrightarrow{a_{n-1}} q_{n-1} \xrightarrow{a_n} q$$

and

$$q \xrightarrow{a_n^{-1}} q_{n-1} \xrightarrow{a_{n-1}^{-1}} \dots \xrightarrow{a_2^{-1}} q_1 \xrightarrow{a_1^{-1}} p.$$

Note that this automaton is still trim and dual. Moreover, if $p = q = q_0$, then we observe the following.

Proposition 1.2. *Let $\mathcal{A} = (Q, q_0, E)$ be a trim dual automaton, let $H = L(\mathcal{A})\rho$ and let w be a non-empty word. If \mathcal{B} is the expansion of \mathcal{A} by (q_0, w, q_0) , then $L(\mathcal{B})\rho$ is the subgroup generated by H and $w\rho$, that is, $L(\mathcal{B})\rho = \langle H, w \rangle$.*

Proof. Let \mathcal{C} be the dual automaton consisting of the state q_0 and the states and edges added to \mathcal{A} in the expansion. It is immediate that $L(\mathcal{C})\rho$ is the subgroup of $F(\mathcal{A})$ generated by $w\rho$.

If $u \in L(\mathcal{B})$, we can factor a path $q_0 \xrightarrow{u} q_0$ according to the successive visits of state q_0 . The resulting factorization of u makes it clear that u is a product of elements of $L(\mathcal{A})$ and $L(\mathcal{C})$. Thus, $L(\mathcal{B})$ is the submonoid generated by $L(\mathcal{A}) \cup L(\mathcal{C})$, and $L(\mathcal{B})\rho$ is the subgroup generated by $L(\mathcal{A})\rho$ and $w\rho$. This concludes the proof. \square

1.2. REDUCED INVERSE AUTOMATA

The automaton \mathcal{A} is called *inverse* if it is deterministic, trim and dual. It is *reduced* if every state $q \in Q$ lies in some path from q_0 to q_0 , labeled by a (possibly empty) reduced word. We note the following result, a cousin of [24, Thm 1.16].

Proposition 1.3. *If \mathcal{A} and \mathcal{B} are reduced inverse automata such that $L(\mathcal{A})\rho = L(\mathcal{B})\rho$, then \mathcal{A} and \mathcal{B} are isomorphic.*

Proof. Let $\mathcal{A} = (Q, q_0, E)$ and $\mathcal{B} = (P, p_0, D)$ be reduced inverse automata such that $L(\mathcal{A})\rho = L(\mathcal{B})\rho$. We construct an isomorphism φ between \mathcal{A} and \mathcal{B} as follows. We first let $\varphi(q_0) = p_0$.

Let $q \in Q$. Since \mathcal{A} is reduced, there exist reduced words u and v such that the word uv is reduced, $q_0 \xrightarrow{u} q$ and $q \xrightarrow{v} q_0$. Then $uv \in L(\mathcal{A})\rho$, so $uv \in L(\mathcal{B})\rho$, and hence $uv \in L(\mathcal{B})$ by Fact 1.1. Thus uv labels a path in \mathcal{B} from p_0 to p_0 , and we let $\varphi(q)$ be the unique state in P such that $p_0 \xrightarrow{u} \varphi(q) \xrightarrow{v} p_0$.


We first verify that φ is well defined. Suppose that uv and $u'v'$ are reduced words such that $q_0 \xrightarrow{u} q \xrightarrow{v} q_0$ and $q_0 \xrightarrow{u'} q \xrightarrow{v'} q_0$ in \mathcal{A} . We want to show that if $p_0 \xrightarrow{u} p \xrightarrow{v} p_0$ and $p_0 \xrightarrow{u'} p' \xrightarrow{v'} p_0$ in \mathcal{B} , then $p = p'$. We note that $u'v$ labels a path from q_0 to q_0 in \mathcal{A} . If $u'v$ is a reduced word, then by the same reasoning as above, $u'v$ labels a path in \mathcal{B} from p_0 to p_0 , say, $p_0 \xrightarrow{u'} p'' \xrightarrow{v} p_0$ and the deterministic property of \mathcal{B} implies that $p' = p'' = p$.

If $u'v$ is not reduced, and a is the first letter of v , then the last letter of u' is a^{-1} while the last letter of u is not a^{-1} . Therefore $u'u^{-1}$ is reduced, $u'u^{-1} \in L(\mathcal{A})$ and again, there is a path in \mathcal{B} of the form $p_0 \xrightarrow{u'} p'' \xrightarrow{u^{-1}} p_0$. By determinism, it follows that $p' = p'' = p$.

This shows that φ is well defined. A dual construction yields a well-defined mapping ψ from P to Q such that, whenever $p_0 \xrightarrow{u} p \xrightarrow{v} p_0$ in \mathcal{B} and uv is a reduced word, then $q_0 \xrightarrow{u} \psi(p) \xrightarrow{v} q_0$ in \mathcal{A} . Using the determinism of \mathcal{A} and \mathcal{B} , it is now immediate that $\psi \circ \varphi$ is the identity on Q and $\varphi \circ \psi$ is the identity on P .

There remains to verify that φ and φ^{-1} are homomorphisms. The case of φ^{-1} is dual of that of φ and we treat only the latter. That is, we want to show that if

(q, a, q') is a transition in \mathcal{A} , then $(\varphi(q), a, \varphi(q'))$ is a transition in \mathcal{B} . Let uv and $u'v'$ be reduced words such that $q_0 \xrightarrow{u} q \xrightarrow{v} q_0$ and $q_0 \xrightarrow{u'} q' \xrightarrow{v'} q_0$. In particular, we have $p_0 \xrightarrow{u} \varphi(q) \xrightarrow{v} p_0$ and $p_0 \xrightarrow{u'} \varphi(q') \xrightarrow{v'} p_0$ in \mathcal{B} .



If uav' is reduced, then in \mathcal{B} , there is a path from p_0 to p_0 labeled uav' , and by determinism, there is a transition $(\varphi(q), a, \varphi(q'))$. If uav' is not reduced, then either ua is not reduced or av' is not reduced. If ua is not reduced, then $u = u_1a^{-1}$ and by determinism, $q_0 \xrightarrow{u_1} q'$. As in the first part of the proof, it follows that at least one of u_1v' and $u_1u'^{-1}$ is reduced, so $p_0 \xrightarrow{u_1} \varphi(q')$ in \mathcal{B} and hence there is a transition $(\varphi(q), a, \varphi(q'))$. The case where av' is not reduced is handled symmetrically, and this concludes the proof. \square

Let H be a subgroup of $F(A)$. Say that an automaton \mathcal{A} on alphabet A represents H if \mathcal{A} is reduced and inverse and if $L(\mathcal{A})\rho = H$. Proposition 1.3 shows that there exists at most one such automaton, and we denote it by $\Gamma_A(H)$ if it exists. We now discuss the existence and the construction of $\Gamma_A(H)$ when H is finitely generated. (As it turns out, $\Gamma_A(H)$ always exists, but our interest in this paper is restricted to the finite rank case.)

Let \mathcal{A} be an automaton and let p, q be distinct states of \mathcal{A} . The automaton obtained from \mathcal{A} by identifying states p and q is constructed as follows: its state set is $Q \setminus \{p, q\} \cup \{n\}$, where n is a new state; its initial state is q_0 (or n if p or q is equal to q_0); and its set of transitions is obtained from E by replacing everywhere p and q by n . If \mathcal{A} is trim or dual, then so is the automaton obtained from \mathcal{A} by identifying a pair of states.

Now let \mathcal{A} be a dual automaton. If \mathcal{A} is not deterministic, there exist transitions (r, a, p) and (r, a, q) with $p \neq q$ and $a \in \tilde{A}$. Identifying p and q yields a new dual automaton \mathcal{B} , and we say that \mathcal{B} is obtained from \mathcal{A} by an *elementary reduction of type 1*.

Fact 1.4. Let \mathcal{A} be a dual automaton and let \mathcal{B} be obtained from \mathcal{A} by an elementary reduction of type 1. Then $L(\mathcal{A})\rho = L(\mathcal{B})\rho$.

Proof. It is easily seen that $L(\mathcal{A}) \subseteq L(\mathcal{B})$. For the converse, we use the notation given above: in \mathcal{B} , the states p and q of \mathcal{A} are replaced with a new state n . Let $u \in L(\mathcal{B})$. Then there exists a path labeled u from the initial state of \mathcal{B} (say, q_0) to itself. If that path does not visit state n , then u also labels a path from q_0 to itself in \mathcal{A} and hence $u \in L(\mathcal{A})$.

If that path does visit state n , we consider the factorization of u given by the passage of that path through n : we have $u = u_0 u_1 \cdots u_r$, $r \geq 1$ and

$$q_0 \xrightarrow{u_0} n \xrightarrow{u_1} n \cdots n \xrightarrow{u_r} q_0.$$

It follows that in \mathcal{A} , u_i -labelled paths exist, with end states p or q (or q_0). Then one of u_0 and $u_0 a^{-1} a$ labels a path in \mathcal{A} from q_0 to q . Similarly, one of u_r and $a^{-1} a u_r$ labels a path from q to q_0 (making due allowance if p or q is equal to q_0). And for each $1 \leq i \leq r$, one of u_i , $a^{-1} a u_i$, $u_i a^{-1} a$ and $a^{-1} a u_i a^{-1} a$ labels a path in \mathcal{A} from q to q . Therefore, there exists a path in \mathcal{A} of the form $q_0 \xrightarrow{v} q_0$ such that $u\rho = v\rho$, which concludes the proof. \square

Now assume that \mathcal{A} is a deterministic dual automaton.

Fact 1.5. Let \mathcal{A} be an inverse automaton. Then \mathcal{A} is non-reduced if and only if there exist states $q \neq q_0$ and p , and a letter $a \in \tilde{A}$ such that the only transitions of \mathcal{A} involving q are (p, a, q) and (q, a^{-1}, p) .

In graph-theoretic terms, this means that \mathcal{A} is reduced if and only if no vertex of \mathcal{A} has degree one (more precisely: no vertex is adjacent to a single A -labeled edge), except possibly q_0 .

Proof. By definition, \mathcal{A} is not reduced if and only if there exists a state q that does not lie on any path from q_0 to itself, labeled by a reduced word. We first observe that the state q cannot be equal to q_0 since the empty word is reduced, and labels a path from q_0 to itself. Suppose now that there exist reduced words u, v with distinct last letters, labeling paths from q_0 to q : then uv^{-1} is a reduced word, labeling a path from q_0 to itself and visiting q . On the other hand, if every reduced word labeling a path from q_0 to q ends with, say, letter $a \in \tilde{A}$, then every path from q_0 to itself visiting q has a non-reduced label. Thus \mathcal{A} is not reduced if and only if there exists a state $q \neq q_0$ and every reduced word labeling a path from q_0 to q ends with the same letter. By determinism, this is equivalent to the existence of another state p such that the transitions involving q are (p, a, q) and (q, a^{-1}, p) . \square

Let \mathcal{A} be inverse and not reduced, and let a, p, q be as in Fact 1.5. If \mathcal{B} is obtained from \mathcal{A} by omitting state q and the transitions involving it, we observe that \mathcal{B} is again an inverse automaton, and we say that \mathcal{B} is obtained from \mathcal{A} by an *elementary reduction of type 2*.

Fact 1.6. Let \mathcal{A} be an inverse automaton and let \mathcal{B} be obtained from \mathcal{A} by an elementary reduction of type 2. Then $L(\mathcal{A})\rho = L(\mathcal{B})\rho$.

Proof. Let a be a letter and let p, q be states of \mathcal{A} as in Fact 1.5, and let us assume that \mathcal{B} is obtained from \mathcal{A} by omitting state q and the transitions involving it. It is easily seen that $L(\mathcal{B}) \subseteq L(\mathcal{A})$. Conversely, let $u \in L(\mathcal{A})$. By Fact 1.1, $u\rho \in L(\mathcal{A})$. Now Fact 1.5 shows that the path $q_0 \xrightarrow{u} \rho q_0$ in \mathcal{A} cannot visit state q , since $u\rho$ is a reduced word. It follows that this path is also a path in \mathcal{B} , $u\rho \in L(\mathcal{B})$ and hence $u\rho \in L(\mathcal{B})\rho$. \square

Let \mathcal{A} be a trim, dual automaton, and let \mathcal{B} be an automaton obtained by iteratively performing elementary reductions, first of type 1 until the automaton is inverse, and then of type 2 until none is possible. Then \mathcal{B} is a reduced inverse automaton, we write $\mathcal{B} = \mathcal{A}\rho$ and we say that \mathcal{B} is obtained from \mathcal{A} by reduction. Moreover, Facts 1.4 and 1.6 show that $L(\mathcal{A})\rho = L(\mathcal{B})\rho$.

This leads directly to the well-known algorithm to construct a reduced inverse automaton representing a given finitely generated subgroup H . Let h_1, \dots, h_n be generators of H , and let us consider the automaton obtained from the trivial automaton (one vertex q_0 , no transitions) by performing successively expansions by (q_0, h_i, q_0) ($1 \leq i \leq n$) and then reducing the automaton. It follows from Propositions 1.2 and 1.3 that the resulting automaton is $\Gamma_{\mathcal{A}}(H)$. Note that it does not matter which set of generators of H was used, nor in which order the elementary reductions were performed.

Remark 1.7. This construction of $\Gamma_{\mathcal{A}}(H)$ is well known, and can be described in many different ways, notably in terms of immersions over the bouquet of circles (Stallings [23]) or of closed inverse submonoids of a free inverse monoid (Margolis and Meakin [13]).

Fact 1.8. There is a well-known converse to the above construction: if \mathcal{A} is a reduced inverse automaton and $H = L(\mathcal{A})\rho$, then H has finite rank and a basis for H can be computed as follows (see Stallings [23]). Given a spanning tree T of the (graph underlying the) automaton \mathcal{A} , for each state p , let u_p be the reduced word labeling a path from q_0 to p inside the tree T . For each transition $e = (p, a, q)$, let $b_e = u_p a u_q^{-1}$: then a basis of H consists of the elements b_e , where e runs over the transitions $e = (p, a, q)$ not in T and such that $a \in A$.

We note that, given a finite set h_1, \dots, h_n of elements of $F(A)$ with total length $\ell = \sum_i |h_i|$, one can construct $\Gamma_{\mathcal{A}}(H)$ in time at most $O(\ell^2)$ and $\Gamma_{\mathcal{A}}(H)$ has $v \leq \ell - n + 1$ states. Moreover, finding a basis of H can be done in time at most $O(v^2)$ ($O(v \log^* v)$) according to Touikan's announcement [25], and the rank of H is equal to $e - v + 1$, where e is the number of edges in $\Gamma_{\mathcal{A}}(H)$.

1.3. ON THE COMPLEXITY OF WHITEHEAD AND OTHER ALGORITHMS

It is well known that one can decide, given a subgroup H of a finite rank free group F , whether H is a free factor of F . We briefly describe here the main known algorithms and discuss their complexity.

Let H be a finitely generated subgroup of a free group F of rank r , with basis A . Let h_1, \dots, h_n be a generating set of H . By the results summarized in Fact 1.8, up to a quadratic time computation, we may assume that h_1, \dots, h_n is a basis of H . Let $\ell = |h_1| + \dots + |h_n|$ be the total length of the tuple $(h_i)_i$, and let $d = r - n$ be the rank difference between F and H – which we assume to be positive, since H can be a proper free factor of F only if $n < r$.

Federer and Jónsson (see [12, Prop. I.2.26]) gave the following observation and decision procedure: H is a free factor of F if and only if there exist d words h_{n+1}, \dots, h_r , each of length at most $\max\{|h_i| \mid 1 \leq i \leq n\}$, such that h_1, \dots, h_r

generate the whole of F . The resulting algorithm requires testing every suitable d -tuple of reduced words on alphabet A . Each of these tests (does a certain r -tuple of words generate F ?) takes time polynomial in the total length of the r -tuple, and hence in $d\ell$. However, the number of tests is $O(r^{d\ell})$, which is exponential in ℓ and d .

This approach leads to the following.

Fact 1.9. Deciding whether $H \leq_{\text{ff}} K$ is in NP , with respect to $d\ell$.

Proof. To verify that $H \leq_{\text{ff}} K$, we need to guess d words of length at most ℓ , and verify that together with H , they generate F , which can be done in $O((d\ell)^2)$. \square

Another approach is based on the use of Whitehead automorphisms. We refer the readers to [12, Sec. I.4] for the definition of these automorphisms, it suffices to note here that the set W of Whitehead automorphisms of F which do not preserve length, has exponential cardinality (in terms of r). A result of Whitehead [12, Prop. I.4.24] shows the following: if there exists an automorphism φ of F such that the total length of $(\varphi(h_i))_i$ is strictly less than ℓ , then there exists such an automorphism in W . In particular, an algorithm to compute the minimum total length of an automorphic image of the tuple $(h_i)_i$ consists in repeatedly applying the following step: try every automorphism $\psi \in W$ until the total length of $(\psi(h_i))_i$ is strictly less than the total length of $(h_i)_i$; if such a ψ exists, replace $(h_i)_i$ by $(\psi(h_i))_i$; otherwise, stop and output the total length of $(h_i)_i$.

This applies to the decision of the free factor relation since $H \leq_{\text{ff}} F$ if and only if there exists an automorphism φ mapping a basis of H to a subset of A . Thus an algorithm consists in first computing a basis of H , and assuming that h_1, \dots, h_n is a basis, verifying whether the minimum total length of $(\varphi(h_i))_i$ when φ runs over the automorphisms of F , is exactly n . This algorithm may require $O((\ell - n)\text{card}(W))$ steps, each of which consists in computing the image of a tuple of length at most ℓ under an automorphism, and hence has complexity $O(\ell)$. Thus the time complexity of this algorithm is $O(\ell^2 \text{card}(W))$, which is quadratic in ℓ and exponential in r .

A variant of this algorithm was established by Gersten [5], who showed that a similar method applies to find the minimum size (number of vertices) of $\Gamma_A(\varphi(H))$, when φ runs over the automorphisms of $F(A)$. It is clear that H is a free factor of $F(A)$ if and only if there exists an automorphism φ such that $\Gamma_A(\varphi(H))$ has a single vertex. The time complexity is computed as above, where the number of vertices of $\Gamma_A(H)$ is substituted for the total length of a basis for H . As noted earlier, this number of vertices is usually substantially smaller than the total length of a basis, but the two values are linearly dependent, so the order of magnitude of the time complexity is not modified, notably the exponential dependence in r .

Remark 1.10. The discussion of Whitehead's algorithm above concerns only the so-called *easy part* of the algorithm (see for instance Kapovich, Miasnikov and Shpilrain [9]). Results by Miasnikov and Shpilrain [15], Khan [10], and most recently by Donghi Lee [11] on the possible polynomial complexity of the *hard part*

of the algorithm also consider the rank of the ambient free group as a constant, and do not question the actual exponential dependence in that parameter.

2. A CAREFUL LOOK AT THE EXPANSIONS AND REDUCTIONS OF INVERSE AUTOMATA

Let \mathcal{A} be a reduced inverse automaton.

Let \mathcal{B} be obtained from \mathcal{A} by performing an expansion, say by (p, w, q) , and then reducing the resulting automaton. In this situation, we write $\mathcal{A} \xrightarrow{\text{exp}}^{(p,w,q)} \mathcal{B}$, or simply $\mathcal{A} \xrightarrow{\text{exp}} \mathcal{B}$. We distinguish two special cases.

- If the reduction following the expansion does not involve identifying or omitting states of \mathcal{A} , or equivalently if \mathcal{A} embeds in \mathcal{B} , we say that \mathcal{B} is obtained from \mathcal{A} by a *reduced expansion* and we write $\mathcal{A} \xrightarrow{\text{re}}^{(p,w,q)} \mathcal{B}$ or $\mathcal{A} \xrightarrow{\text{re}} \mathcal{B}$.

- If the states p and q are equal to the distinguished state q_0 of \mathcal{A} , we say that \mathcal{B} is obtained from \mathcal{A} by an *e-step* and we write $\mathcal{A} \xrightarrow{\text{e}}^w \mathcal{B}$, or simply $\mathcal{A} \xrightarrow{\text{e}} \mathcal{B}$.

Finally, let \mathcal{B} be obtained from \mathcal{A} by identifying two distinct vertices p and q , and then reducing the resulting automaton. Then we say that \mathcal{B} is obtained from \mathcal{A} by an *i-step* and we write $\mathcal{A} \xrightarrow{\text{i}}^{p=q} \mathcal{B}$, or simply $\mathcal{A} \xrightarrow{\text{i}} \mathcal{B}$.

Note that if $\mathcal{A} \xrightarrow{\text{exp}} \mathcal{B}$, $\mathcal{A} \xrightarrow{\text{re}} \mathcal{B}$, $\mathcal{A} \xrightarrow{\text{e}} \mathcal{B}$ or $\mathcal{A} \xrightarrow{\text{i}} \mathcal{B}$, then \mathcal{B} is a reduced inverse automaton.

We first record a few facts.

Fact 2.1. Let u be a reduced word labeling a path in \mathcal{A} from a state p to a state p' , and from a state q to a state q' ,

$$p \xrightarrow{u} p', \quad q \xrightarrow{u} q'.$$

By definition of the reduction of dual automata, the identification of p and q implies that of p' and q' , and the converse holds as well. Thus $\mathcal{A} \xrightarrow{\text{i}}^{p=q} \mathcal{B}$ if and only if $\mathcal{A} \xrightarrow{\text{i}}^{p'=q'} \mathcal{B}$.

Let us now examine in detail the effect of an operation of the form $\xrightarrow{\text{exp}}$.

Fact 2.2. Let p, q be states of \mathcal{A} and let w be a non-empty reduced word. Let u be the longest prefix of w that can be read in \mathcal{A} from state p , and let v be the longest suffix of w that can be read in \mathcal{A} to state q (that is, v^{-1} is the longest prefix of w^{-1} that can be read in \mathcal{A} from state q). We distinguish two cases:

- (1) If $|u| + |v| < |w|$, then $w = uw'v$ for some non-empty reduced word w' . If we let p' (resp. q') be the end (resp. start) state of the path labeled u (resp. v) and starting in p (resp. ending in q),

$$p \xrightarrow{u} p' \xrightarrow{w'} q' \xrightarrow{v} q,$$

then the reduction process on the result of the expansion of \mathcal{A} by (p, w, q) identifies the $|u|$ first edges and the $|v|$ last edges of the added path with

existing edges of \mathcal{A} , so that $\mathcal{A} \xrightarrow{\text{exp}}^{(p,w,q)} \mathcal{B}$ if and only if $\mathcal{A} \xrightarrow{\text{exp}}^{(p',w',q')} \mathcal{B}$ and the latter is a reduced expansion.

- (2) If $|u| + |v| \geq |w|$, then there exist words x, y, z , possibly empty, such that $u = xy$, $v = yz$ and $w = xyz$. Let p', p'', q', q'' be the states of \mathcal{A} defined by the following paths

$$p \xrightarrow{x} p' \xrightarrow{y} p'', \quad q' \xrightarrow{y} q'' \xrightarrow{z} q.$$

Then $\mathcal{A} \xrightarrow{\text{exp}}^{(p,w,q)} \mathcal{B}$ if and only if $\mathcal{A} \xrightarrow{i}^{p'=q'} \mathcal{B}$, if and only if $\mathcal{A} \xrightarrow{i}^{p''=q''} \mathcal{B}$.

We derive from Fact 2.2 the following statement.

Proposition 2.3. *Let \mathcal{A} and \mathcal{B} be inverse automata. If $\mathcal{A} \xrightarrow{e}^w \mathcal{B}$, then $\mathcal{A} \xrightarrow{i}$ \mathcal{B} or $\mathcal{A} \xrightarrow{\text{re}}^{(p,u,q)} \mathcal{B}$ for some states p and q and a reduced word u such that $|u| \leq |w|$.*

The following converse statements are derived from Facts 2.1 and 2.2.

Proposition 2.4. *Let \mathcal{A} be a reduced inverse automaton, let $H = L(\mathcal{A})\rho$, let u and v be reduced words labeling paths $q \xrightarrow{v} q_0 \xrightarrow{u} p$ in \mathcal{A} , and suppose that $\mathcal{A} \xrightarrow{i}^{p=q} \mathcal{B}$. Then $\mathcal{A} \xrightarrow{e}^{uv} \mathcal{B}$ and $L(\mathcal{B})\rho = \langle H, uv \rangle$.*

In particular, $\text{rank}(L(\mathcal{B})\rho) \leq 1 + \text{rank}(H)$. If $\text{rank}(L(\mathcal{B})\rho) = 1 + \text{rank}(H)$, then $H \leq_{\text{ff}} L(\mathcal{B})\rho$ and if C is a basis of H , then $C \cup \{uv\}$ is a basis of $L(\mathcal{B})\rho$.

Proof. Let \mathcal{A}' be the expansion of \mathcal{A} by (q_0, uv, q_0) . The analysis in Fact 2.2 (2) shows that a step in the reduction of \mathcal{A}' is provided by the automaton obtained in identifying p and q . By Proposition 1.2 we have $L(\mathcal{B})\rho = \langle H, uv \rangle$ (hence the rank inequality), and the uniqueness statement in Proposition 1.3 then shows that $\mathcal{A} \xrightarrow{e}^{uv} \mathcal{B}$.

Let us now assume that $\text{rank}(L(\mathcal{B})\rho) = 1 + \text{rank}(H)$, and that C is a basis of H . Then $C \cup \{uv\}$ is a generating set of $L(\mathcal{B})\rho$ with cardinality equal to the rank of $L(\mathcal{B})\rho$, so $C \cup \{uv\}$ is a basis of that subgroup by [12, Prop. I.3.5]. \square

Proposition 2.5. *Let \mathcal{A} and \mathcal{B} be reduced inverse automata, let w be a reduced word such that $\mathcal{A} \xrightarrow{\text{re}}^{(p,w,q)} \mathcal{B}$, let $H = L(\mathcal{A})\rho$, and let u and v be reduced words labeling paths $q \xrightarrow{v} q_0 \xrightarrow{u} p$ in \mathcal{A} . Then $\mathcal{A} \xrightarrow{e}^{uvw} \mathcal{B}$ and $L(\mathcal{B})\rho = \langle H, uvw \rangle$.*

Moreover, $H \leq_{\text{ff}} L(\mathcal{B})\rho$, $\text{rank}(L(\mathcal{B})\rho) = 1 + \text{rank}(H)$ and if C is a basis of H , then $C \cup \{uvw\}$ is a basis of $L(\mathcal{B})\rho$.

Proof. Since the expansion of \mathcal{A} by (p, w, q) is a reduced expansion, the word uvw is reduced and the expansion by (q_0, uvw, q_0) falls in the situation described in Fact 2.2 (1). Together with Proposition 1.2, it follows that $\mathcal{A} \xrightarrow{e}^{uvw} \mathcal{B}$, which concludes the verification that of the first part of the proposition.

The free factor and the rank statements follow from the statement on a basis for $L(\mathcal{B})\rho$, which is a direct consequence of the definition of a reduced expansion and of the discussion on bases and spanning trees in Fact 1.8. \square

We now introduce a measure of the *cost* of a reduced expansion or an i-step σ , written $\lambda(\sigma)$: if σ is an i-step, then $\lambda(\sigma) = 0$; if σ is a reduced expansion,

$\sigma \xrightarrow{\text{re}}_{(p,w,q)}$, its cost is the length of w , $\lambda(\sigma) = |w|$. We extend this notion of cost to finite sequences of i -steps and reduced expansions: if $\bar{\sigma} = (\sigma_1, \dots, \sigma_n)$ is such a sequence, we let

$$\lambda(\bar{\sigma}) = (\lambda(\sigma_1), \dots, \lambda(\sigma_n)).$$

Finally, we introduce an order relation on the set of finite sequences of non-negative integers. Let $\bar{k} = (k_1, \dots, k_n)$ and $\bar{\ell} = (\ell_1, \dots, \ell_m)$ be such sequences. We say that $\bar{k} \preceq \bar{\ell}$ if

- either $n < m$,
- or $n = m$ and $\sum_{i=1}^n k_i < \sum_{i=1}^m \ell_i$,
- or $n = m$, $\sum_{i=1}^n k_i = \sum_{i=1}^m \ell_i$ and \bar{k} precedes $\bar{\ell}$ in the lexicographic order.

It is routine to check that \preceq is a well-order on the set of finite sequences of non-negative integers, which is stable under the concatenation of sequences. We write $\bar{k} \prec \bar{\ell}$ if $\bar{k} \preceq \bar{\ell}$ and $\bar{k} \neq \bar{\ell}$.

Proposition 2.6. *Let \mathcal{A} , \mathcal{A}' and \mathcal{B} be inverse automata such that \mathcal{A}' is obtained from \mathcal{A} by a reduced expansion σ_1 and \mathcal{B} is obtained from \mathcal{A}' by an i -step σ_2 ,*

$$\mathcal{A} \xrightarrow{\text{re}} \mathcal{A}' \xrightarrow{i} \mathcal{B}.$$

Then there exist a sequence of reduced expansions or i -steps $\bar{\sigma}'$ of length 1 or 2 such that \mathcal{B} is obtained from \mathcal{A} by applying the steps in $\bar{\sigma}'$ and $\lambda(\bar{\sigma}') \prec \lambda(\sigma_1, \sigma_2)$.

Proof. Suppose that $\mathcal{A} \xrightarrow{\text{re}}_{(p,w,q)} \mathcal{A}' \xrightarrow{i}^{r=s} \mathcal{B}$ with $r \neq s$. The cost of this sequence of transformations is $(|w|, 0)$.

Let Q be the state set of \mathcal{A} and let u and v be reduced paths,

$$q \xrightarrow{v} q_0 \xrightarrow{u} p.$$

Then uvw is a reduced word and $L(\mathcal{A}')\rho = \langle L(\mathcal{A})\rho, uvw \rangle$ by Proposition 2.5. We distinguish three cases, depending whether or not r and s lie in Q .

Case 1: Both r and s are in Q . Let x and y be reduced words labeling paths in \mathcal{A} $s \xrightarrow{y} q_0 \xrightarrow{x} r$. Then the same words label similar paths in \mathcal{A}' and it follows from Proposition 2.4 that

$$L(\mathcal{B})\rho = \langle L(\mathcal{A}')\rho, xy \rangle = \langle L(\mathcal{A})\rho, uvw, xy \rangle.$$

Let also \mathcal{A}'' and \mathcal{B}' be determined by $\mathcal{A} \xrightarrow{i}^{r=s} \mathcal{A}'' \xrightarrow{e}^{uvw} \mathcal{B}'$. Then $L(\mathcal{B}')\rho$ is also equal to $\langle L(\mathcal{A})\rho, xy, uvw \rangle$, so that $\mathcal{B} = \mathcal{B}'$ by Proposition 1.3.

Note that the words u and v label paths from and into state q_0 in \mathcal{A}'' as well. It follows from Proposition 2.3 that, if $uvw \notin L(\mathcal{A}'')$, then \mathcal{B} can be obtained from \mathcal{A}'' by an i -step or by a reduced expansion of the form $\xrightarrow{\text{re}}_{(t,z,t')}$ with $|z| \leq |w|$.

Thus \mathcal{B} is obtained from \mathcal{A} either by a sequence of 1 or 2 transformations, of cost 0 or $(0, k)$ with $0 \leq k \leq |w|$. This is \prec -less than $(|w|, 0)$, as expected.

Case 2: Exactly one of r and s is in Q . Without loss of generality, we may assume that $r \in Q$ and $s \notin Q$. Let z be a reduced word labeling a path from q_0 to r in \mathcal{A} , and hence also in \mathcal{A}' . Let g be the unique reduced word labeling a path from p to s in \mathcal{A}' , using only edges that were not in \mathcal{A} . By assumption, g is a proper, non-empty prefix of w . Moreover, by Propositions 1.2 and 2.4,

$$L(\mathcal{B})\rho = \langle L(\mathcal{A}')\rho, ugz^{-1} \rangle = \langle L(\mathcal{A})\rho, uvw, ugz^{-1} \rangle.$$

Let h be the longest common suffix of g and z , so that $g = g'h$, $z = z'h$, $g'z'^{-1}$ is reduced and we have the following paths in \mathcal{A}' ,

$$q_0 \xrightarrow{z'} r' \xrightarrow{h} r, \quad p \xrightarrow{g'} s' \xrightarrow{h} s.$$

Fact 2.1 shows that $\mathcal{A}' \xrightarrow{i}^{r'=s'} \mathcal{B}$, so we may assume that $h = 1$, $g = g'$ and $z = z'$. There is a possibility that the word g is now empty (if h was in fact equal to g), but in that case, we are returned to the situation of Case 1, with $s' = p$. Thus we may still assume that $g \neq 1$. In particular, the word ugz^{-1} is reduced.

Then let \mathcal{A}'' and \mathcal{B}' be defined by $\mathcal{A} \xrightarrow{e}^{ugz^{-1}} \mathcal{A}'' \xrightarrow{e}^{uvw} \mathcal{B}'$. Again $L(\mathcal{B}')\rho = \langle L(\mathcal{A})\rho, uvw, ugz^{-1} \rangle$, so $\mathcal{B} = \mathcal{B}'$ by Proposition 1.3.

Proposition 2.3 states that each e-step can be replaced by an i-step or by a reduced expansion of cost bounded above by the cost of the e-step. Going back to Fact 2.2, we see that the e-step $\mathcal{A} \xrightarrow{e}^{ugz^{-1}} \mathcal{A}''$ can be replaced by a transformation of cost $k \leq |g|$ since both u and z can be read from state q_0 in \mathcal{A} (in fact, of cost exactly $|g|$ by definition of g). As for the e-step $\mathcal{A}'' \xrightarrow{e}^{uvw} \mathcal{B}'$, it can be replaced by a transformation of cost $\ell \leq |w| - |g|$ since ug (a prefix of uw) and v can be read to state q_0 in \mathcal{A}'' .

Now, it suffices to verify that $(k, \ell) \prec (|w|, 0)$, which is easily done if we observe that $k + \ell \leq |w|$ and $k < |w|$.

Case 3: Neither r nor s is in Q . Without loss of generality, we may assume that r occurs before s along the w -labeled path from p to q . Thus, the word w factors as $w = w_1w_2w_3$ and the path in \mathcal{A}' made of edges added to \mathcal{A} factors as

$$p \xrightarrow{w_1} r \xrightarrow{w_2} s \xrightarrow{w_3} q.$$

Since $r \neq s$ and these vertices are not in Q , each of the three factors w_1, w_2, w_3 is non-empty. Moreover,

$$L(\mathcal{B})\rho = \langle L(\mathcal{A}')\rho, ww_1w_3v \rangle = \langle L(\mathcal{A})\rho, uvw, ww_1w_3v \rangle.$$

Let h be the longest common suffix of w_1 and w_3^{-1} , so that $w_1 = w'_1h$, $w_3 = h^{-1}w'_3$, $w'_1w'_3$ is reduced and we have the following paths in \mathcal{A}' ,

$$p \xrightarrow{w'_1} r' \xrightarrow{h} r \xrightarrow{w_2} s \xleftarrow{h} s' \xrightarrow{w'_3} q.$$

Proposition 2.1 shows that $\mathcal{A}' \xrightarrow{i}^{r'=s'} \mathcal{B}$, so we may assume that $h = 1$, $w_1 = w'_1$ and $w_3 = w'_3$. There is a possibility that the words w_1 or w_3 be now empty (if h

was in fact equal to w_1 or w_3), but in that case, we are returned to the situation of Cases 1 or 2, with $r' = p$ or $s' = q$. Thus we may still assume that $w_1 \neq 1$ and $w_3 \neq 1$. In particular, the word uw_1w_3v is reduced.

Then let \mathcal{A}'' and \mathcal{B}' be defined by $\mathcal{A} \xrightarrow{uw_1w_3v} \mathcal{A}'' \xrightarrow{uvw} \mathcal{B}'$. Then $L(\mathcal{B}')\rho = \langle L(\mathcal{A})\rho, uvw, uw_1w_3v \rangle$, so $\mathcal{B} = \mathcal{B}'$ by Proposition 1.3.

As in Case 2, we use Fact 2.2 to verify that the e-step $\mathcal{A} \xrightarrow{uw_1w_3v} \mathcal{A}''$ can be replaced by a reduced expansion of cost $k = |w_1w_3|$ since u and v are the maximal prefix and suffix of uw_1w_3v that can be read from and to state q_0 in \mathcal{A} . As for the e-step $\mathcal{A}'' \xrightarrow{uvw} \mathcal{B}$, it can be replaced by a reduced expansion of cost $\ell = |w_2|$ since uw_1 and w_3v are the maximal prefix and suffix of uvw that can be read from and to state q_0 in \mathcal{A}'' .

Now, it suffices to verify that $(k, \ell) \prec (|w|, 0)$, which is easily done if we observe that $k + \ell = |w|$ and $k < |w|$. \square

3. DECIDING THE FREE FACTOR RELATION

3.1. A GEOMETRIC CHARACTERIZATION OF FREE FACTORS

We put together the technical results from Section 2 to prove the following characterization of free factors.

Theorem 3.1. *Let H, K be finitely generated subgroups of $F = F(A)$ and assume that $d = \text{rank}(K) - \text{rank}(H) > 0$. Then H is a free factor of K if and only if the inverse automaton $\Gamma_A(H)$ can be transformed in $\Gamma_A(K)$ by a sequence of $d' \leq d$ i-steps followed by $d - d'$ reduced expansions.*

Proof. We first observe that H is a free factor of K if and only if there exist d elements k_1, \dots, k_d of $F(A)$ such that $\langle H \cup \{k_1, \dots, k_d\} \rangle = K$. This follows from the fact that an r -element generating set in a rank r free group, is a basis [12, Prop. I.3.5].

By definition of e-steps, this means that $H \leq_{\text{ff}} K$ if and only if $\Gamma_A(H)$ yields $\Gamma_A(K)$ by a sequence of d e-steps.

Now Propositions 2.3, 2.4 and 2.5 show that this is equivalent to the fact that $\Gamma_A(H)$ yields $\Gamma_A(K)$ by a sequence of d i-steps or reduced expansions.

Since \preceq is a well-order on the set of finite sequences of non-negative integers, we may consider a sequence $\bar{\sigma}$ of d i-steps and reduced expansions leading from $\Gamma_A(H)$ to $\Gamma_A(K)$, which is \preceq -minimal. Proposition 2.6, together with the stability of \preceq under concatenation, then shows that the i-steps in $\bar{\sigma}$ come before the reduced expansions. Thus, $H \leq_{\text{ff}} K$ if and only if $\Gamma_A(H)$ yields $\Gamma_A(K)$ by a sequence of d' i-steps followed by $d - d'$ reduced expansions. \square

It follows from the discussion on bases and spanning trees in Fact 1.8 that if $\Gamma_A(H)$ embeds in $\Gamma_A(K)$, then H is a free factor of K . Not every free factor of K occurs that way, and those that do are called the *graphical free factors of K (with respect to A)*. It is easily verified that H is a graphical free factor of K if and only if $\Gamma_A(H)$ yields $\Gamma_A(K)$ by a sequence of reduced expansions.

Corollary 3.2. *Let H, K be finitely generated subgroups of $F = F(A)$ and assume that $d = \text{rank}(K) - \text{rank}(H) > 0$. Then H is a free factor of K if and only if the inverse automaton $\Gamma_A(H)$ can be transformed by a sequence of $d' \leq d$ i-steps into some $\Gamma_A(L)$ such that $\text{rank}(L) = \text{rank}(H) + d'$ and L is a graphical free factor of K with respect to A .*

Proof. Let us first assume that $H \leq_{\text{ff}} K$. By Theorem 3.1, for some $d' \leq d$, $\Gamma_A(H)$ can be taken to some $\Gamma_A(L)$ by a sequence of d' i-steps, and $\Gamma_A(L)$ can be taken to $\Gamma_A(K)$ by a sequence of $d - d'$ re-steps. Since $\text{rank}(K) = \text{rank}(H) + d$ and an i-step or a re-step can increment the rank by at most one (Propositions 2.4 and 2.5), L must have rank exactly $\text{rank}(H) + d'$.

Conversely, suppose that a sequence of $d' \leq d$ i-steps takes $\Gamma_A(H)$ to $\Gamma_A(L)$ in such a way that $\text{rank}(L) = \text{rank}(H) + d'$ and $\Gamma_A(L)$ embeds in $\Gamma_A(K)$. By Propositions 2.4 and 2.5 again, we have $H \leq_{\text{ff}} L \leq_{\text{ff}} K$, and hence $H \leq_{\text{ff}} K$. \square

Remark 3.3. We observe the following by-product of the proof of Corollary 3.2: if $\Gamma_A(H)$ can be transformed into $\Gamma_A(L)$ by a sequence of d' i-steps such that $\text{rank}(L) = d' + \text{rank}(H)$, then for every i-step $\mathcal{A} \rightarrow_i \mathcal{B}$ occurring in that sequence, we have $\text{rank}(L(\mathcal{B})\rho) = 1 + \text{rank}(L(\mathcal{A})\rho)$. We say that such an i-step is *rank-incrementing*.

In the special case where $K = F$, we have the following statement.

Corollary 3.4. *Let H be a finitely generated subgroup of $F = F(A)$, let A_0 be the set of letters in A that occur in $\Gamma_A(H)$ and let $d = |A_0| - \text{rank}(H) = \text{rank}(F(A_0)) - \text{rank}(H)$. Then H is a free factor of F if and only if $d \geq 0$ and $\Gamma_A(H)$ can be transformed into a one-vertex automaton by a sequence of d i-steps.*

Proof. By Corollary 3.2, H is a free factor of $F(A)$ if and only if $\Gamma_A(H)$ yields a one-vertex automaton \mathcal{B} by a sequence of $d' \leq |A| - \text{rank}(H)$ i-steps, in such a way that $L(\mathcal{B})\rho$ has rank $d' + \text{rank}(H)$. Now the set of letters occurring in such an automaton \mathcal{B} must be exactly A_0 , so $L(\mathcal{B})\rho = F(A_0)$ and $d' = d$. \square

3.2. DECIDING WHETHER $H \leq_{\text{ff}} F$

We first consider the problem of deciding whether a given finitely generated subgroup H of $F = F(A)$ is a free factor of F . With the notation of Corollary 3.4 and Remark 3.3, the algorithm to decide whether $H \leq_{\text{ff}} F$ consists of the following. We need to explore the sequences of rank-incrementing i-steps, starting from $\Gamma_A(H)$ and of length $d = |A_0| - \text{rank}(H)$. Then $H \leq_{\text{ff}} F$ if and only if one of the automata occurring at the end of one of these sequences has a single vertex. Note that no automaton obtained after less than d i-steps could be a single-vertex automaton.

These automata can be viewed as nodes of a tree, rooted at $\Gamma_A(H)$, in which the children of a node are the automata produced by a rank-incrementing i-step. Let n be the number of states of $\Gamma_A(H)$. Then every automaton occurring along this tree has at most n states.

If \mathcal{A} is such an automaton, then \mathcal{A} has at most $\frac{1}{2}(n^2 - n)$ pairs of distinct vertices, and hence at most $\frac{1}{2}(n^2 - n)$ children, each of which has at most $n - 1$ states. The computation of the children of \mathcal{A} is done by computing all the (at most $\frac{1}{2}(n^2 - n)$) possible i-steps, computing the ranks of the corresponding subgroups and retaining only those of rank $1 + \text{rank}(L(\mathcal{A})\rho)$. It follows from Fact 1.8 that the cost of the computation of the children of \mathcal{A} is $O(n^4)$.

Moreover, there are at most $O(n^{2d-2})$ nodes of the tree at depth at most $d - 1$, and the computation of these nodes and their children requires time $O(n^{2d+2})$. For each of the $O(n^{2d})$ automata at depth d , the verification whether the automaton has a single node takes constant time, so the total cost of the algorithm is $O(n^{2d+2})$.

Finally, if H is given by a finite set of generators, of total length ℓ , we recall that computing $\Gamma_A(H)$ takes time $O(\ell^2)$ and that $\Gamma_A(H)$ has at most ℓ states and ℓ edges. This discussion justifies the following statement.

Theorem 3.5. *Let A be a fixed alphabet. Then there is an algorithm which, given reduced words h_1, \dots, h_n in $F(A)$ with total length ℓ , decides whether the subgroup H generated by the h_i is a free factor of $F(A)$ in time $O(\ell^{2d+2})$, where $d = |A_0| - \text{rank}(H)$ and A_0 is the set of letters in A that occur in the h_i .*

Remark 3.6. The tree exploration described above can be speeded up by the following observation. If the automaton \mathcal{A} occurs in a sequence of rank-incrementing i-steps from $\Gamma_A(H)$ to the one-vertex automaton $\Gamma_A(F(A_0))$ (a *winning sequence*), then $L(\mathcal{A})\rho \leq_{\text{ff}} F(A_0)$, which implies that $L(\mathcal{A})\rho$ is a free factor of every subgroup of $F(A_0)$ containing it. It follows that, if some i-step $\mathcal{A} \rightarrow_i \mathcal{B}$ does not increment the rank, then \mathcal{A} does not occur in a winning sequence, that is, we may ignore the subtree below \mathcal{A} .

There are undoubtedly other implementation tricks and ideas that can reduce the decision process, however without changing the worst-case complexity.

Remark 3.7. In view of Touikan's announcement (see Fact 1.8), the time complexity in the above theorem can be lowered to $O(\ell^{2d+1} \log^* \ell)$.

The above discussion of complexity depends on a *uniform cost assumption*, by which we assume that the elementary operations on A (reading or writing a letter, comparing two letters) require unit time. In other words, we have been assuming that A is fixed, and not part of the input. We will shortly consider the problem of deciding whether $H \leq_{\text{ff}} K$ where both H and K are part of the input, and for the purpose of that discussion, we reconsider the algorithm and the complexity established above under the *bit cost assumption*: we let r be the cardinality of A , and we consider that each letter is identified by a bit string of length at most $\log r$, so that the elementary operations on A require $O(\log r)$ units of time.

Let $f(\ell, r)$ be the complexity of reducing a dual automaton on \tilde{A} with at most ℓ vertices and ℓ edges. The computation of the automaton obtained by an i-step from an ℓ -vertex automaton takes time at most $f(\ell, r)$. To compute the tree of rank-incrementing i-steps rooted at $\Gamma_A(H)$, we need to compute the at most $O(\ell^2)$ children of at most ℓ^{2d-2} nodes, which requires time $O(\ell^{2d} f(\ell, r))$. Finally,

deciding whether an automaton has a single vertex takes constant time, so the total complexity of the algorithm is $O(\ell^{2d}f(\ell, r))$.

We now give an upper bound of $f(\ell, r)$. Let \mathcal{A} be an A -automaton with at most ℓ states and ℓ edges. The identifiers of states require space $O(\log \ell)$ and the identifiers of letters require space $O(\log r)$. Moreover, we assume that the set of states and the set \tilde{A} are linearly ordered, and equipped with a constant time next function.

The automaton can be viewed as the lexicographically ordered list of all triples (u, a, v) such that either $a \in A$ and there is an a -labeled edge from state u to state v , or $\bar{a} \in A$ and there is a \bar{a} -labeled edge from state v to state u . Each entry of this list requires space $O(\log \ell + \log r) = O(\log(\ell r))$, and the list contains at most ℓ entries. In particular, a complete scan of the list takes time $O(\ell \log(\ell r))$.

The reduction of \mathcal{A} consists in performing elementary reductions of type 1 as long as it is possible, and then elementary reductions of type 2. To find out whether a type 1 reduction is possible, one needs to scan the list to find two (consecutive) cells of the form (u, a, v) and (u, a, v') , requiring $O(\ell \log(\ell r))$. Performing the identification consists in reading through the list, replacing every occurrence of v' by v , and reordering the list: this takes time $O(\ell \log(\ell r))$. To find out whether an elementary reduction of type 2 is possible, one scans the list to find a vertex u such that there is a unique triple of the form (u, a, v) in the list: again, this takes time $O(\ell \log(\ell r))$. Finally, performing the reduction consists simply in deleting the entries (u, a, v) and (v, \bar{a}, u) in the list.

Summarizing, since we will perform at most ℓ reductions, we can take the function $f(\ell, r)$ to be equal to $\ell^2 \log(\ell r)$. This yields the following statement.

Theorem 3.8. *There is an algorithm which, given an alphabet A of size r and reduced words $h_1, \dots, h_n \in F(A)$ with total length ℓ , decides whether the subgroup H generated by the h_i is a free factor of $F(A)$ in time $O(\ell^{2d+2} \log(\ell r))$, where $d = |A_0| - \text{rank}(H)$ and A_0 is the set of letters in A that occur in the h_i .*

3.3. DECIDING WHETHER $H \leq_{\text{ff}} K$

We now suppose that H and K are subgroups of $F(A)$, given by sets of generators with total length ℓ , and we consider the problem of deciding whether $H \leq_{\text{ff}} K$. The algorithm is the following: we first compute $\Gamma_A(K)$ and we verify whether each generator of H can be read as a loop at q_0 , the designated vertex of $\Gamma_A(K)$. If not, then H is not contained in K , and cannot be a free factor.

We now assume that $H \leq K$ and we compute a spanning tree T of the graph $\Gamma_A(K)$. This spanning tree determines a basis B of K , indexed by the edges of $\Gamma_A(K)$ that are not in T (see Fact 1.8), and we express the given generators of H in terms of this basis: it suffices to read the generators of H as loops at the designated vertex in $\Gamma_A(K)$, and to record the sequence of edges traversed and not in T . In particular, each generator of H is expressed as a reduced word in \tilde{B}^* , that is no longer than its expression as a reduced word in \tilde{A}^* . We then use this expression of the generators of H to construct $\Gamma_B(H)$, and the algorithm in

Section 3.2 to decide whether H is a free factor of the ambient free group, namely K .

Let us now discuss the complexity of this algorithm, assuming that H and K are given by tuples of generators of total length ℓ and that F has rank r . As discussed in the previous section, computing $\Gamma_A(K)$ takes time $O(\ell^2 \log(\ell r))$. Running a word $w \in F(A)$ in the automaton $\Gamma_A(K)$ requires reading sequentially each letter of w (time $O(\log r)$) and looking for the corresponding transition in the table representing the automaton (time $O(\ell \log(\ell r))$). If w has length ℓ' , this takes time $O(\ell' \ell \log(\ell r))$. In order to decide whether H is contained in K , this must be done for every generator of H , and since the total length of these generators is at most ℓ , this requires time $O(\ell^2 \log(\ell r))$.

Assuming that H is indeed contained in K , the next step is to construct a spanning tree T of $\Gamma_A(H)$, for instance by marking certain edges in the list representing $\Gamma_A(K)$. Again, this can be done in time $O(\ell^2 \log(\ell r))$. The positively labeled edges of $\Gamma_A(K)$ not in T are in bijection with a basis B of K . In particular, the rank of K is at most ℓ , an upper bound of the number of edges in $\Gamma_A(K)$. The elements of B , seen as words in $F(A)$, consist of the label of a path in T , followed by the label of an edge not in T , followed by the label of a path in T . In particular, their length is at most twice the number of vertices of $\Gamma_A(K)$ plus one, that is $O(\ell)$. But we do not need to compute these words: it suffices to number (from 1 to $\text{rank}(K) \leq \ell$) the positively labeled edges not in T . Then, for each generator h of H , reading h in $\Gamma_A(K)$ from state q_0 and keeping track of the (identifier of the) edges traversed that are not in T , provides an expression of h in B . Moreover, the total length of the expression of the generators of H in this basis of K is at most ℓ .

We now simply apply Theorem 3.8 to a length ℓ set of generators of a subgroup of a free group of rank at most ℓ , to justify the following corollary.

Corollary 3.9. *Given tuples of generators for subgroups H and K of $F(A)$, with total length ℓ , one can decide whether the subgroup H is a free factor of K in time $O(\ell^{2d+2} \log(\ell r))$, where $d = \text{rank}(K) - \text{rank}(H)$.*

Remark 3.10. Corollary 3.2 suggests an alternative algorithm to decide whether $H \leq_{\text{ff}} K$: one may explore the sequences of rank-incrementing i-steps of length at most $d = \text{rank}(K) - \text{rank}(H)$, starting from $\Gamma_A(H)$ and producing only representations of subgroups contained in K . Each $\Gamma_A(L)$ occurring in such a sequence must be checked to verify whether it can be embedded in $\Gamma_A(K)$: if it can, then $H \leq_{\text{ff}} K$, and if it cannot, then the automata produced by a rank-incrementing i-step from it must be computed and checked in their turn.

3.4. COMPLEMENT OF A FREE FACTOR

By definition of a free factor (see the introduction), $H \leq_{\text{ff}} K$ if there exists a basis C of H and a disjoint set $D \subseteq K$ such that $C \cup D$ is a basis of K . In that case, the subgroup generated by D is called a *complement of H in K* . It is easy to see that this complement is not uniquely defined.

If H is indeed a free factor of the free group F , the algorithm described in Section 3.2 also provides a sequence of rank-incrementing i -steps taking $\Gamma_A(H)$ to $\Gamma_A(F(A_0))$, where A_0 is the set of letters of A that occur in the generators of H . Repeated applications of Proposition 2.4 then yield a basis of $F(A_0)$ containing a basis of H , and hence a complement of H in $F(A_0)$.

We can be a little more precise. If \mathcal{A} is a reduced inverse automaton, let $\text{diam}_{q_0}(\mathcal{A})$, the q_0 -diameter of \mathcal{A} , be the longest shortest path in \mathcal{A} from q_0 to a state, that is

$$\text{diam}_{q_0}(\mathcal{A}) = \max_q \left(\min\{|u| \mid u \text{ is a reduced word and } q_0 \xrightarrow{u} q \text{ in } \mathcal{A}\} \right).$$

Then Proposition 2.4 shows that if $\mathcal{A} \rightarrow_i \mathcal{B}$ and $L(\mathcal{A})\rho \leq_{\text{ff}} L(\mathcal{B})\rho$, then a complement of $L(\mathcal{A})\rho$ is generated by a word of length at most $2 \text{diam}_{q_0}(\mathcal{A})$.

This justifies the following statement.

Proposition 3.11. *Let H be a free factor of $F(A)$. Then one can construct effectively a basis for a complement of H in $F(A)$, consisting of words of length at most $2 \text{diam}_{q_0}(\Gamma_A(H))$.*

It is interesting to observe that this is a graphical analogue – and a minor improvement – of Federer and Jónsson’s result (see [12, Prop. I.2.26]) mentioned in Section 1.3. Indeed, if h_1, \dots, h_n are a set of generators of H , then every edge of $\Gamma_A(H)$ is on a loop at q_0 labeled by some h_i , and hence $\max_i |h_i| \geq 2 \text{diam}_{q_0}(\Gamma_A(H))$.

Let us now consider the problem of constructing a basis of a complement of H in K , where H and K are given finitely generated subgroups of $F(A)$ and $H \leq_{\text{ff}} K$. The algorithm described in Remark 3.10 (based on Corollary 3.2) shows that one can construct effectively a graphical free factor L of K such that $H \leq_{\text{ff}} L$. As above, in view of Proposition 2.4, the same algorithm can produce a basis of a complement of H in L consisting of words of length at most $2 \text{diam}_{q_0}(\Gamma_A(H))$. There remains to construct a basis for a complement of L in K .

In view of Fact 1.8 (and using the notation therein), one can find a (basis of a) complement of L in K by considering a spanning tree of $\Gamma_A(L)$, extending it to a spanning tree T of $\Gamma_A(K)$, and considering the elements b_e associated with edges of $\Gamma_A(K)$ that are neither in $\Gamma_A(L)$ nor in T . In particular, the words b_e have length at most $1 + 2 \text{diam}_{q_0}(\Gamma(K))$. Thus we have the following statement.

Proposition 3.12. *Let H, K be finitely generated subgroups of $F(A)$. If $H \leq_{\text{ff}} K$, then one can construct effectively a basis for a complement of H in K , consisting of words of length at most $\max(2 \text{diam}_{q_0}(\Gamma_A(H)), 1 + 2 \text{diam}_{q_0}(\Gamma_A(K)))$.*

The authors wish to thank both anonymous referees for their many helpful comments, which have substantially contributed to the improvement of the paper. All remaining shortcomings are naturally ours.

REFERENCES

- [1] J. Almeida. *Finite semigroups and universal algebra*, World Scientific Publishing, Singapore, 1994.
- [2] I. Anshel, M. Anshel, B. Fisher, D. Goldfeld. New key agreement protocols in braid group cryptography. In: *CT-RSA 2001*, Lect. Notes in Computer Science **2020**, Springer (2001), 1-15.
- [3] S. Cleary, J. Taback. Metric properties of the lamplighter group as an automata group, *Contemp. Math. Series.* **372**, AMS, 2005
- [4] P. Dehornoy. Braid-based cryptography, *Contemporary Mathematics* **360** (2004) 5-33
- [5] S. Gersten. On Whitehead's algorithm, *Bull. Am. Math. Soc.* **10** (1984) 281-284.
- [6] K. Henckell, S.W. Margolis, J.-E. Pin, J. Rhodes. Ash's type II theorem, profinite topology and Malcev products, *Int. J. Algebra and Comput.* **1** (1991) 411-436.
- [7] M. Kambites, P. V. Silva, B. Steinberg. The spectra of lamplighter groups and Cayley machines, *Geometriae Dedicata* (to appear).
- [8] I. Kapovich, A.G. Miasnikov. Stallings foldings and subgroups of free groups, *J. Algebra*, **248**, 2 (2002), 608-668.
- [9] I. Kapovich, P. Schupp, V. Shpilrain. Generic properties of Whitehead's algorithm and isomorphism rigidity of random one-relator groups, *Pacific J. Math.* **223** (2006) 113-140.
- [10] B. Khan. The structure of automorphic conjugacy in the free group of rank two. In: *Proc. Special Session on Interactions between Logic, Group Theory and Computer Science*, Contemp. Mathematics **349** (2004).
- [11] D. Lee. A tighter bound for the number of words of minimum length in an automorphic orbit, eprint [arXiv:math.GR/0401269](https://arxiv.org/abs/math/0401269), 2006.
- [12] R. Lyndon, P. Schupp. *Combinatorial group theory*, Springer, (1977, reprinted 2001).
- [13] S. Margolis, J. Meakin. Free inverse monoids and graph immersions, *Intern. J. Algebra and Computation* **3** (1993) 79-100.
- [14] S. Margolis, M. Sapir, P. Weil. Closed subgroups in pro- \mathbf{V} topologies and the extension problem for inverse automata, *Intern. J. Algebra and Computation* **11** (2001) 405-445.
- [15] A.G. Miasnikov, V. Shpilrain. Automorphic orbits in free groups, *J. Algebra* **269** (2003) 18-27.
- [16] A.G. Miasnikov, E. Ventura, P. Weil. Algebraic extensions in free groups, in preparation.
- [17] D. Perrin. Automata, in (J. Leeuwen ed.) *Handbook of Theoretical Computer Science*, vol. B, Elsevier, 1990.
- [18] J.-E. Pin, *Variétés de langages formels*, Masson, Paris (1984); English translation: *Varieties of formal languages*, Plenum, New-York (1986).
- [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] V.M. Sidelnikov, M.A. Cherepnev, V.Y. Yaschenko. Systems of open distribution of keys on the basis of non-commutative semigroups, *Russ. Acad. Nauk Dokl.* **332-5** (1993). English translation: *Russian Acad. Sci. Dokl. Math.* **48-2** (1994) 383-386.
- [22] P. V. Silva, B. Steinberg. On a class of automata groups generalizing lamplighter groups, *Intern. J. Algebra and Computation* **15** (2005) 1213-1234.
- [23] J. Stallings. Topology of finite graphs, *Inventiones Mathematicae* **71** (1983) 551-565.
- [24] J. Stephen. *Applications of automata theory to presentations of monoids and inverse monoids*, Ph.D. Dissertation, University of Nebraska, 1987.
- [25] N. Touikan. A fast algorithm for Stallings's folding process, preprint, 2005, www.math.mcgill.ca/~touikan/crypto_seminar/FastFolding.pdf.
- [26] E. Ventura. On fixed subgroups of maximal rank, *Comm. Algebra*, **25** (1997), 3361-3375.

Communicated by (The editor will be set by the publisher).
(The dates will be set by the publisher).