Computing closures of finitely generated subgroups of the free group

Pascal Weil[1]

The aim of this paper is to present computing techniques for the finitely generated subgroups of the free group. The focus is the computation of the closures of a finitely generated subgroup for certain profinite topologies, namely the pro-p and the pro-nilpotent topology. However the first section concentrates on the classical manipulation of finitely generated subgroups when more elementary operations are concerned: computing the rank, the index or a basis; calculating an intersection.

The emphasis in the paper is on the algorithmic techniques and one will find no proof here. We purposely restricted ourselves to presenting only the necessary concepts, reserving as much space as possible for examples meant to illustrate the computing techniques we describe. Bibliographical pointers are given to papers where proofs and justifications can be found.

As mentioned above, Section 1 does not deal with topological properties. It provides the tools and algorithms to perform basic operations on finitely generated subgroups of the free group. The central tool is the bijective correspondence between these subgroups and certain finite structures, the so-called reduced inverse automata. The idea behind this association goes back to the early part of the 20th century [7, Chapter 11] and has been developed in Stallings [8] in an algebraic topology terminology, and more or less in the terminology adopted here in the more recent papers [4, 1, 5], where proofs of the statements of the section can be found. Of particular interest is a simple proof of Hall's extension theorem, which states that every finitely generated subgroup of the free group is a free factor of a finite index subgroup [2].

[1]LaBRI, CNRS – 351 cours de la Libération – 33405 Talence Cedex, France
Pascal.Weil@labri.u-bordeaux.fr http://dept-info.labri.u-bordeaux.fr/~weil

In Section 2, we consider the pro-$\mathbf{V}$ topology on the free group, where $\mathbf{V}$ is a pseudovariety of finite groups (e.g. the finite p-groups, or the finite nilpotent groups). These topologies were introduced by Hall [2], who proved that when $\mathbf{V}$ is the class of all finite groups, then the finite index subgroups are exactly the open subgroups, and that every finitely generated subgroup is closed. The proof of this last result is a variant of the proof of Hall's extension theorem.

In [6], Ribes and Zalesskiĭ showed that the pro-p closure of a finitely generated subgroup is finitely generated, and is effectively computable. In [5], Margolis, Sapir and Weil reworked Ribes and Zalesskiĭ's algorithm to show that the computation can be performed in polynomial time. In the process, they clarified the relation between the profinite topologies and the extension property examplified by Hall's extension theorem. These results are presented in Section 2.

The actual algorithm to compute the pro-p closure is presented in Section 3, as well as one to compute the pro-nilpotent closure. Details on these algorithms can be found in [5]. We conclude with a brief discussion of an open problem: the computation of the pro-solvable closure of a finitely generated subgroup of the free group.

1 Subgroups of the free group and automata

Throughout the paper, A is a fixed finite, non-empty set. Let $\bar{A}$ be a disjoint copy of A. For each word u in $(A \cup \bar{A})^*$, the free monoid on $A \cup \bar{A}$, we denote by $\gamma(u)$ the unique reduced word obtained from u by iteratively deleting consecutions of letters of the form $a\bar{a}$ or $\bar{a}a$ ($a \in A$).

Let $F(A)$ be the free group on A. It is well-known that $F(A)$ can be identified with the set of reduced words on alphabet $A \cup \bar{A}$, where the product of reduced words u and v is $\gamma(uv)$. That is, the mapping $\gamma\colon (A \cup \bar{A})^* \to F(A)$ is an onto morphism and its restriction to the reduced words is a bijection.

1.1 Inverse automata

The basic tool to represent subgroups of the free group is borrowed from automata theory. Proofs of the facts stated in this section can be found in [5].

1.1.1 Basic definition

The automata we will be using in this paper are the so-called *inverse automata*. An inverse automaton on alphabet A is a triple (Q, i, δ) where Q is the finite non-empty set of states, $i \in Q$ is the distinguished initial-terminal state and $\delta: Q \times (A \cup \bar{A})$ is a mapping such that, for each $a \in A$, the mappings $q \mapsto \delta(q, a)$ and $q \mapsto \delta(q, \bar{a})$ are mutually reciprocal partial one-to-one transformations of Q. We also require an accessibility condition which we explain below.

The *runs* of the automaton (Q, i, δ) are defined inductively as follows. Let $p, q \in Q$. The automaton has a run from p to q on the empty word if and only if $p = q$. If $a \in A \cup \bar{A}$, the automaton has a run on a from p to q if $\delta(p, a) = q$. Finally, if $u = va$ with $v \in (A \cup \bar{A})^*$ and $a \in A \cup \bar{A}$, then the automaton has a run from p to q on u if there exists a state r such that the automaton has a run on v from p to r, and such that $q \in \delta(r, a)$.

This can be expressed in terms of an extension of the mapping δ to $Q \times (A \cup \bar{A})^*$: we let $\delta(p, 1) = \{p\}$; and if $v \in (A \cup \bar{A})^*$ and $a \in A \cup \bar{A}$, we let $\delta(p, va) = \bigcup_{r \in \delta(p,v)} \delta(r, a)$. Then the automaton has a run from p to q on u if and only if $q \in \delta(p, u)$.

The accessibility condition which is part of the definition of an inverse automaton is the following:

(Accessibility) there exists a run from i to each state $q \in Q$.

The language *accepted* by the automaton $\mathcal{A} = (Q, i, \delta)$ is the set $L(\mathcal{A})$ of words $u \in (A \cup \bar{A})^*$ such that $\mathcal{A}$ has a run on u from i to i.

We represent inverse automata on A as A-labeled graphs with a distinguished vertex, as in the following example. The accessibility condition above corresponds to the connectedness of the graph.

Example 1.1 Here is a 4-state inverse automaton on alphabet $A = \{a, b\}$.

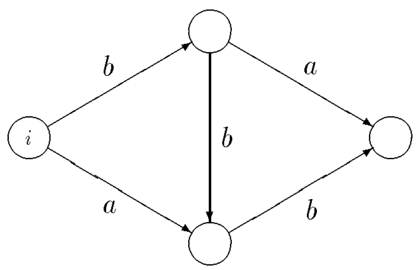

The vertices of the graph are the states of the automaton, and the edges labeled a or b represent the transitions of the automaton with the same labels. The runs of the automaton from state p to q are the labels of paths in this graph, from p to q, provided we allow traversing an a-labeled (resp. b-labeled) edge in the reverse direction and reading a letter $\bar{a}$ (resp. $\bar{b}$).

Because of this graphic representation of inverse automata, we often call the vertices states, and the A-labeled transitions edges of the automaton.

1.1.2 Reduced automata

We say that the automaton (Q, i, δ) is *reduced* if every state $q \neq i$ is visited along a run from i to i labeled by a reduced word. This is equivalent to the following condition:

$$\forall\ q \neq i,\ \exists\ a \neq b \in A \cup \bar{A}\ \text{ such that } \delta(q, a) \text{ and } \delta(q, b) \text{ both exist.}$$

That is, every vertex but possibly the initial-terminal vertex i is adjacent to 2 distinct positively labeled edges. Said again differently, the graph representing the automaton does not have vertices of valence 1, except perhaps for the initial-terminal vertex.

Example 1.2 The inverse automaton of Example 1.1 is reduced. The first of the following inverse automata is not reduced. The second one is.

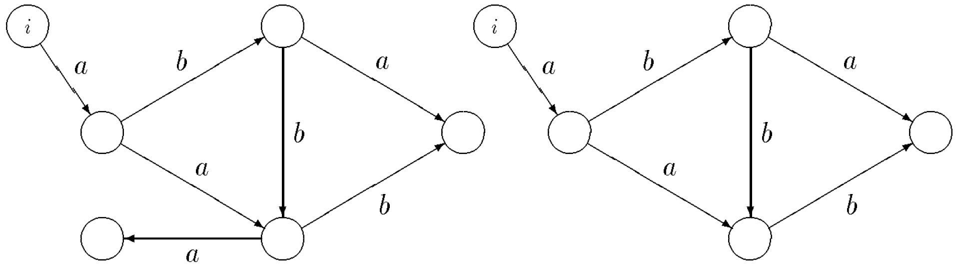

Given an inverse automaton (Q, i, δ), we may consider the sub-automaton (Q_r, i, δ_r) where Q_r consists of state i and all the states which are visited along some run from i to i on a reduced word, and where δ_r is the restriction of δ to a map from $Q_r \times A$ into Q_r. It is immediate that this automaton is inverse and reduced: it is called the *reduced core* of the automaton (Q, i, δ).

There is a procedure to reduce a finite inverse automaton to its reduced core: it suffices to iteratively remove a state of valence 1 (other than the initial-terminal state) and all the transitions involving that state. It is easily verified that this process terminates and that its final product is the reduced core of the given automaton.

Example 1.3 The second automaton in Example 1.2 is the reduced core of the first one.

The above procedure to transform an inverse automaton into its reduced core is analogous to the procedure to reduce a word on alphabet $A \cup \bar{A}$. It may be puzzling that, in the process of reducing an inverse automaton, the initial-terminal vertex is allowed to remain with valence 1. Here the analogy is with the notion of a cyclically reduced word, that is, a reduced word u which cannot be factored as $u = v^{-1}wv$ for a non-empty word v. Following this analogy, we call *cyclically reduced core* of $\mathcal{A}(H)$ the A-labeled graph (without a distinguished vertex) obtained from $\mathcal{A}(H)$ by iteratively removing the vertices of valence 1, including possibly the initial-terminal vertex.

1.1.3 Morphisms and congruences

Let $\mathcal{A} = (Q, i, \delta)$ and $\mathcal{B} = (P, j, \varepsilon)$ be inverse automata over the same alphabet. A *morphism* from $\mathcal{A}$ to $\mathcal{B}$ is a mapping $\kappa: Q \to P$ such that

- $\kappa(i) = j$
- if $\delta(q, a)$ exists, then $\varepsilon(\kappa(q), a)$ exists and $\kappa(\delta(q, a)) = \varepsilon(\kappa(q), a)$.

One can verify that if such a morphism exists, then it is unique [5, Prop. 2.6]. We say that κ is an isomorphism if the mapping κ is a bijection and κ^{-1} is a morphism as well. Note that it does not suffice that κ is a bijection, if $\mathcal{A}$ is a copy of $\mathcal{B}$ where some edges have been deleted.

As usual, there is a dual notion of congruence on an inverse automaton. A *congruence* on an inverse automaton (Q, i, δ) is an equivalence relation $\sim$ on Q such that if $p, q \in Q$ and $p \sim q$, if $a \in A \cup \bar{A}$ and both $\delta(p, a)$ and $\delta(q, a)$ exist, then $\delta(p, a) \sim \delta(q, a)$. If $\sim$ is a congruence on $\mathcal{A} = (Q, i, \delta)$, the quotient set $Q/\sim$ may be equipped with an inverse automaton structure as follows: $\mathcal{A}/\sim = (Q/\sim, [i], \delta_\sim)$ with $\delta_\sim([q], a) = [\delta(p, a)]$ for any state p $\sim$-equivalent to q such that $\delta(p, a)$ exists. If no such state exists, then $\delta_\sim([q], a)$ is undefined. (Here $[q]$ denotes the $\sim$-class of state q.)

Of course, the natural quotient map is a morphism from $\mathcal{A}$ to $\mathcal{A}/\sim$. Conversely, if κ is a morphism from $\mathcal{A}$ to $\mathcal{B}$, then κ induces a congruence on $\mathcal{A}$: two states of $\mathcal{A}$ are $\sim$-equivalent if they have the same image under κ. If π denotes the quotient morphism from $\mathcal{A}$ to $\mathcal{A}/\sim$, then $\kappa = \lambda \circ \pi$ where λ is a morphism from $\mathcal{A}/\sim$ to $\mathcal{B}$ which is injective. That is, $\mathcal{A}/\sim$ is a copy of $\mathcal{B}$ where some edges and vertices have been deleted.

1.2 The correspondence between subgroups and automata

1.2.1 Subgroup associated with an automaton

With each inverse automaton $\mathcal{A} = (Q, i, \delta)$, we associate a subgroup $\mathcal{H}(\mathcal{A})$ of $F(A)$, namely $\mathcal{H}(\mathcal{A}) = \gamma(L(\mathcal{A}))$. That is, $\mathcal{H}(\mathcal{A})$ is the set of words of the form $\gamma(u)$ such that $\delta(i, u) = i$. One can show that it is also the set of reduced words in $L(\mathcal{A})$.

Example 1.4 If $\mathcal{A}$ is as in Example 1.1, then $\mathcal{H}(\mathcal{A}) = \langle aba^{-1}b^{-1}, b^2a^{-1}\rangle$.

It is immediately clear from the figure in Example 1.1 that $aba^{-1}b^{-1}$ and b^2a^{-1} lie in $\mathcal{H}(\mathcal{A})$. However it may not be as clear why these two elements suffice to generate the whole subgroup $\mathcal{H}(\mathcal{A})$. Techniques to perform effectively such computations will be described in Section 1.3.2

One last result of interest on this correspondence is the following.

Proposition 1.5 *Let $\mathcal{A}$ and $\mathcal{B}$ be reduced inverse automata. If $\mathcal{H}(\mathcal{A}) = \mathcal{H}(\mathcal{B})$, then $\mathcal{A}$ and $\mathcal{B}$ are isomorphic.*

1.2.2 Automaton associated with a subgroup

In the converse direction, each finitely generated subgroup H of $F(A)$ is of the form $\mathcal{H}(\mathcal{A})$ for some reduced inverse automaton, which is unique by Proposition 1.5 and is denoted by $\mathcal{A}(H)$.

The existence of such an automaton is proved constructively. Let Y be a finite set of elements of $F(A)$ and let H be the subgroup of $F(A)$ generated by Y. From Y we construct the A-labeled graph $\mathcal{A}(H)$ in three steps.

First, we construct a set of $|Y|$ loops around a common distinguished vertex i, each labeled by an element of Y, with the following convention: an inverse letter a^{-1} ($a \in A$) in a word of Y gives rise to an a-labeled edge in the reverse direction on the corresponding loop.

Then, we iteratively identify identically-labeled pairs of edges starting or ending at the same vertex. (One can show that the order in which these identifications are performed is irrelevant.) The resulting graph $\mathcal{A}(H)$ is a reduced inverse automaton whose associated subgroup of the free group is H. By Proposition 1.5, $\mathcal{A}(H)$ depends on the subgroup H only, and not on the generating set Y.

Example 1.6 Let $Y = \{b^3a^{-1}b^{-1}, aba^{-1}ba^{-1}, b^2a^{-1}\}$. Here are some steps of the computation.

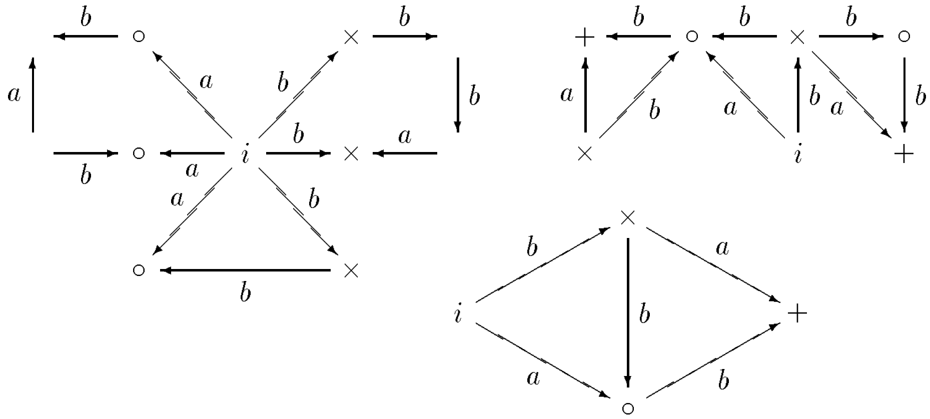

1.3 Operations on subgroups and invariants

1.3.1 Intersection

Let H and K be finitely generated subgroups of the free group. Let us consider the *direct product* of the automata $\mathcal{A}(H)$ and $\mathcal{A}(K)$: it is the A-labeled graph whose vertices are pairs of vertices of $\mathcal{A}(H)$ and $\mathcal{A}(K)$, and such that there is an a-labeled edge from (p, q) to (p', q') if and only if there is one from p to p' in $\mathcal{A}(H)$, and one from q to q' in $\mathcal{A}(K)$. One verifies that if we let (i, i) be the initial-terminal vertex, the connected component of (i, i) is an inverse automaton $\mathcal{A}$ such that $\mathcal{H}(\mathcal{A}) = H \cap K$ (see [8]). Howson's theorem [3] follows.

1.3.2 Rank, basis and index

Let $\mathcal{A} = (Q, i, \delta)$ be an inverse automaton, and let T be a spanning tree of $\mathcal{A}$, that is, a subgraph of $\mathcal{A}$ which uses all the vertices in Q, which satisfies the accessibility condition, and such that any run from a state to itself is labeled by a word u such that $\gamma(u) = 1$ (or equivalently, whose reduced core has one vertex and no edge). For each state $q \in Q$, we let u_q be a reduced word labeling a path in T from i to q (when T is fixed, there exists a unique such word for each q). Let $p_j \xrightarrow{a_j} q_j$ $(1 \leq j \leq k)$ be the A-labeled edges of $\mathcal{A}$ which are not in T. For each j, let $y_j = u_{p_j} a_j u_{q_j}^{-1}$, and let $H = \langle y_1, \ldots, y_k \rangle$. Then $\{y_1, \ldots, y_k\}$ is a basis for H and $\mathcal{A} = \mathcal{A}(H)$. In particular, as T has necessarily $|Q| - 1$ A-labeled edges, the rank of H is equal to $|E| - |Q| + 1$ where E is the set of A-labeled edges of $\mathcal{A}$.

Example 1.7 Let $H = \langle b^3a^{-1}b^{-1}, aba^{-1}ba^{-1}, b^2a^{-1} \rangle$ as in Example 1.6. If we let T consist of the edges along the paths labeled ba and a starting from

i, we get the following basis for H: $\{aba^{-1}b^{-1}, b^2a^{-1}\}$. This justifies the assertion made in Example 1.4. It also shows that the set of generators for H which we considered in Example 1.6 is not a basis for H as it has 3 elements.

The index of H can also be read easily from $\mathcal{A}(H)$. One shows that H has finite index if and only if $\mathcal{A}(H)$ is a permutation automaton, that is, each letter $a \in A$ induces a permutation of Q. In that case, the index of H is equal to $|Q|$. The (finite) group generated by the permutations of Q induced by the letters is called the group associated with H, and it is denoted by $\mathcal{G}(H)$. This group is in fact equal to the quotient of $F(A)$ by $\bigcap_{g\in F(A)} g^{-1}Hg$, the largest normal subgroup of $F(A)$ contained in H.

1.3.3 Overgroups and free factors

Let H and K be finitely generated subgroups of the free group. One can show that $H \subseteq K$ if and only if there exists a morphism from $\mathcal{A}(H)$ into $\mathcal{A}(K)$. As discussed in Section 1.1.3, this morphism is then uniquely determined. It induces a congruence on $\mathcal{A}(H)$, denoted by $\sim_{H,K}$, which can be described as follows: for each state q of $\mathcal{A}(H)$, we fix a reduced word u_q such that $\delta(i, u_q) = q$ in $\mathcal{A}(H)$; then $p \sim_{H,K} q$ if and only if $u_p u_q^{-1} \in K$.

We saw in Section 1.2.1 that if the morphism from $\mathcal{A}(H)$ to $\mathcal{A}(K)$ is an isomorphism, then $H = K$. Two interesting situations arise when that morphism is injective or surjective on vertices.

Let $\kappa\colon \mathcal{A}(H) \to \mathcal{A}(K)$ be a morphism. If κ is injective on the states of $\mathcal{A}(H)$, it follows from the construction in Section 1.3.2 that there exists a basis of H which can be extended to a basis of K. That is, H is a free factor of K. Note however that there are only finitely many subgroups H such that there is an injective morphism from $\mathcal{A}(H)$ to $\mathcal{A}(K)$, whereas K may have infinitely many free factors. Therefore not all free factors of K arise from taking sub-automata of $\mathcal{A}(K)$.

Now assume that the morphism κ is surjective on the set of states. Then we call K an *overgroup* of H. This definition shows that a given subgroup has only finitely many overgroups. Overgroups will play a crucial role in the next sections.

1.3.4 Hall's extension theorem

The following is a well-known theorem of Hall.

Theorem 1.8 *Every finitely generated subgroup of the free group is a free factor of a finite index subgroup.*

We can give a simple and constructive proof of this result. Let H be a finitely generated subgroup of $F(A)$. In the automaton $\mathcal{A}(H)$, every letter of A induces a partial one-to-one transformation of the set of states, Q. Now let us add new edges to $\mathcal{A}(H)$ to complete these partial maps into permutations of Q. This yields a new automaton $\mathcal{B}$ and the subgroup $K = \mathcal{H}(\mathcal{B})$ has finite index and admits H as a free factor (using the results of the previous sections).

Example 1.9 Let $H = \langle aba^{-1}b^{-1}, b^2a^{-1}\rangle$ be as in Example 1.1. The completion process described above can be performed as follows.

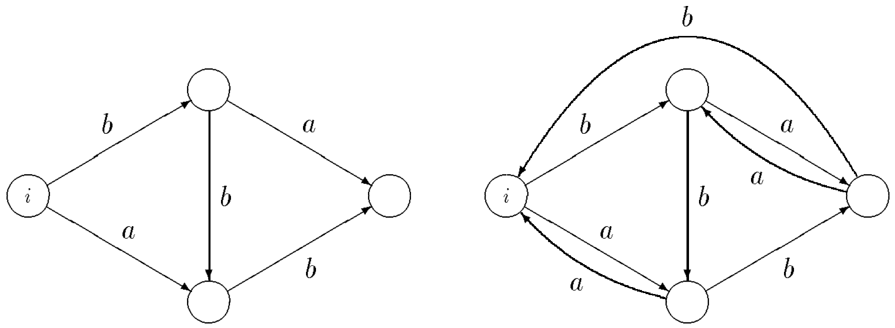

If K is the subgroup associated with the second automaton, then $K = \langle aba^{-1}b^{-1}, b^2a^{-1}, a^2, bab, ba^2b^{-1}\rangle$ has index 4.

2 Profinite topologies and the extension property

Here we examine the interplay between certain topologies that are classical on the free group, and a combinatorial property directly inspired by Hall's extension theorem.

In this section, $\mathbf{V}$ denotes a fixed pseudovariety of groups, that is, a class of finite groups which is closed under taking subgroups, homomorphic images and finite direct products. We also assume $\mathbf{V}$ to be non-trivial, that is, to contain at least one non-trivial group. Important examples which will be discussed in this and the next section are $\mathbf{G}$, the pseudovariety of all finite groups, $\mathbf{G}_p$, the pseudovariety of finite p-groups (where p is a prime), and $\mathbf{G}_{nil}$ and $\mathbf{G}_{sol}$, the pseudovarieties of finite nilpotent and solvable groups. In certain statements, we will also assume that $\mathbf{V}$ is *extension-closed*, that is, if G is a finite group and K is a normal subgroup of G such that K and G/K lie in $\mathbf{V}$, then $G \in \mathbf{V}$.

Whereas the pro-$\mathbf{V}$ topologies (see Section 2.2) go back to Hall [2], the notion of preclosed subgroups was first defined in [5]. The reader is referred to that article for proofs, even though some of the basic results belong to Hall [2] (Proposition 2.5, Section 2.3.1), or to Ribes and Zalesskiĭ [6] (the main part of Proposition 2.9).

2.1 Pre-closed subgroups

We have seen that when H has finite index, the automaton $\mathcal{A}(H)$ is a permutation automaton, which defines the associated finite group $\mathcal{G}(H)$. Among the finite index subgroups of $F(A)$, we distinguish those subgroups H such that $\mathcal{G}(H) \in \mathbf{V}$. We will see (Section 2.2) that these are the open subgroups for a certain topology on $F(A)$, so we will call them $\mathbf{V}$-*open*, or simply *open* subgroups.

Example 2.1 It is not difficult to see that the finite index group K in Example 1.9 has associated group $\mathcal{G}(K) = \mathbb{Z}/4\mathbb{Z}$. Thus, K is $\mathbf{V}$-open when $\mathbf{V} = \mathbf{G}_2$ or $\mathbf{V} = \mathbf{G}_{sol}$, but not when $\mathbf{V} = \mathbf{G}_3$.

We say that a subgroup H is $\mathbf{V}$-*preclosed*, or simply *preclosed* (again, the reason for this terminology will appear more clearly in the sequel) if there exists a $\mathbf{V}$-open subgroup K containing H such that the natural morphism from $\mathcal{A}(H)$ into $\mathcal{A}(K)$ is one-to-one.

In other words, H is preclosed if $\mathcal{A}(H)$ can be extended into a permutation automaton with associated group in $\mathbf{V}$ by adding a finite number of states and transitions. In [5], $\mathbf{V}$-preclosed subgroups are called $\mathbf{V}$-extendible.

Example 2.2 Example 2.1 shows that the subgroup H of Example 1.9 is $\mathbf{V}$-preclosed when $\mathbf{V} = \mathbf{G}_2$ or $\mathbf{V} = \mathbf{G}_{sol}$, in fact whenever $\mathbf{V}$ contains $\mathbb{Z}/4\mathbb{Z}$. We will see however that $\mathbf{V}$ is also $\mathbf{V}$-preclosed when $\mathbf{V} = \mathbf{G}_3$ (Example 3.3) and for every $\mathbf{G}_p$ (p prime).

Since every finite index subgroup is finitely generated, every preclosed subgroup is also finitely generated. Hall's extension theorem, Theorem 1.8 above, shows that when $\mathbf{V} = \mathbf{G}$, the converse holds and every finitely generated subgroup is preclosed.

2.1.1 Preclosed subgroups and free factors of an open subgroup

It follows from the definition that a preclosed subgroup is a free factor of an open subgroup. The converse holds under certain assumptions on $\mathbf{V}$ (Section 2.3.3), but it does not hold in general.

Example 2.3 Let H and K be the following subgroups of the free group.

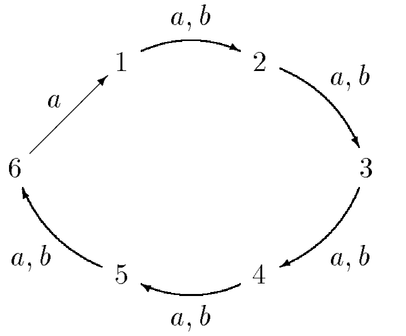

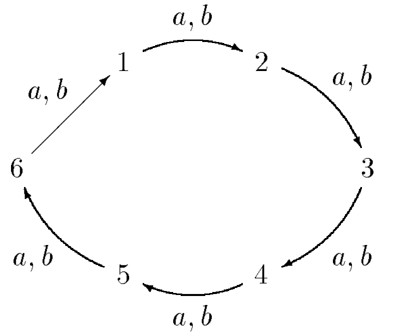

$$H = \langle ab^{-1}, a^2b^{-1}a^{-1}, a^3b^{-1}a^{-2} \quad a^4b^{-1}a^{-3}, a^5b^{-1}a^{-4}, a^6\rangle \qquad K = \langle ab^{-1}, a^2b^{-1}a^{-1}, a^3b^{-1}a^{-2} \quad a^4b^{-1}a^{-3}, a^5b^{-1}a^{-4}, a^6, a^5b\rangle$$

Since $\mathcal{A}(H)$ visibly extends to $\mathcal{A}(K)$ and since K has finite index and $\mathcal{G}(K) = \mathbb{Z}/6\mathbb{Z}$, K is $\mathbf{G}_{nil}$-closed and H is $\mathbf{G}_{nil}$-preclosed.

Observe that $bKb^{-1} = K$. On the other hand, $bHb^{-1} \neq H$, and $\mathcal{A}(bHb^{-1})$ is the following automaton.

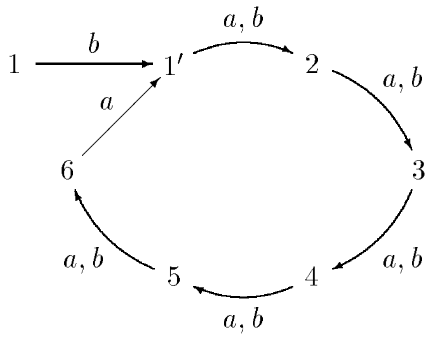

Since H is a free factor of K, bHb^{-1} is also a free factor of $bKb^{-1} = K$. However we will see in Example 3.7 that bHb^{-1} is not $\mathbf{G}_{nil}$-preclosed.

Example 2.3 shows that the notion of preclosed subgroup is not an algebraic notion, since it is not preserved under the automorphisms of the free group.

2.1.2 Preclosure of a finitely generated subgroup

Let H be a finitely generated subgroup of the free group. For every open subgroup K containing H, we can consider the congruence $\sim_{H,K}$ induced on $\mathcal{A}(H)$. The open subgroups K containing H may be infinitely many, but the induced congruences are finitely many, as $\mathcal{A}(H)$ is finite. Now let $\equiv$ be the intersection of these congruences, and let $\tilde{H} = \mathcal{H}(\mathcal{A}(H)/\equiv)$.

By definition, $\tilde{H}$ is an overgroup of H, which we call the *preclosure* of H. One can show the following properties of $\tilde{H}$.

Proposition 2.4 *Let H be a finitely generated subgroup of the free group. There exists an open subgroup K such that $\mathcal{A}(\tilde{H})$ embeds in $\mathcal{A}(K)$, and $\tilde{H}$ is the least preclosed subgroup containing H.*

2.2 Pro-V topologies

If $x, y \in F(A)$ and V is a finite group, we say that V *separates* x and y if there is a morphism $\varphi: F(A) \to V$ such that $\varphi(x) \neq \varphi(y)$.

Now we let $r(x, y)$ be the least order of a group $V \in \mathbf{V}$ such that V separates x and y, and let $d(x, y) = 2^{-r(x,y)}$. One verifies that d is an ultrametric écart, a distance if and only if the free group is residually $\mathbf{V}$. The topology it defines on $F(A)$ is called the *pro-$\mathbf{V}$ topology*. It can be shown to be the initial topology of the morphisms from $F(A)$ into the elements of $\mathbf{V}$.

The operations of $F(A)$ (multiplication and the taking of inverse) are easily seen to be continuous, and $(F(A), d)$ is a precompact topological group (not necessarily Hausdorff). Moreover, every morphism between free groups is uniformly continuous. Open subgroups are characterized as follows.

Proposition 2.5 *Let H be a subgroup of the free group. Then H is open if and only if it has finite index and $\mathcal{G}(H) \in \mathbf{V}$. Moreover, the closure of H is the intersection of the open subgroups containing H.*

Remark. This characterization of open subgroups justifies the terminology adopted in the previous section, before any topology was introduced.

Let H be a subgroup of the free group $F(A)$. It is well-known that H is free too, so for a given $\mathbf{V}$, two different topologies may be considered on H: the restriction to H of the pro-$\mathbf{V}$ topology on $F(A)$, and the pro-$\mathbf{V}$ topology on the free group H itself. It is easily verified that the former is contained in the latter. For the converse, we can note the following result.

Proposition 2.6 *Let H be a subgroup of $F(A)$.*

- *If H is a free factor of $F(A)$, then H is closed and the restriction of the pro-$\mathbf{V}$ topology of $F(A)$ to H, coincides with the pro-$\mathbf{V}$ topology of H.*
- *If $\mathbf{V}$ is extension-closed and H is $\mathbf{V}$-open, then the restriction of the pro-$\mathbf{V}$ topology of $F(A)$ to H, coincides with the pro-$\mathbf{V}$ topology of H.*

Example 2.7 The assumption that $\mathbf{V}$ is extension-closed is necessary in the second statement of Proposition 2.6. This is illustrated by the subgroup bHb^{-1} of Example 2.3: this subgroup is a free factor of a $\mathbf{G}_{nil}$-open subgroup, but it is not $\mathbf{G}_{nil}$-preclosed. We will see in Example 3.7 that in fact H is not $\mathbf{G}_{nil}$-closed, and hence neither is bHb^{-1}.

2.3 Pre-closed vs. closed

There are analogies between the notions of pre-closed and closed subgroups: the closure of H, written $Cl(H)$, is the intersection of the open subgroups K containing H, whereas $\tilde{H}$ is defined by a congruence on $\mathcal{A}(H)$ which is the intersection of the congruences $\sim_{H,K}$ for the same open subgroups K. We will attempt to make this connection clearer in this section.

2.3.1 The case of the pro-G topology

A simple situation prevails when $\mathbf{V} = \mathbf{G}$, the pseudovariety of all finite groups. We already observed that Hall's extension theorem implies that every finitely generated subgroup of the free group is preclosed. In fact, the proof of Hall's theorem can be slightly modified to show that every finitely generated subgroup of the free group is closed. Indeed, in view of Proposition 2.5, we need to show that if H is a finitely generated subgroup of $F(A)$ and $x \notin H$, then there exists an open subgroup K containing H and avoiding x. Let $\mathcal{A}(H) = (Q, i, \delta)$. If $\delta(i, x)$ is not defined, we extend $\mathcal{A}(H)$ by gluing a path labeled x starting from i and remaining in $\mathcal{A}(H)$ as long as possible, to form a new automaton $\mathcal{B}$. If $\delta(i, x)$ is defined, we let $\mathcal{B} = \mathcal{A}(H)$. For instance, here H is as in Example 1.1 and $x = a^2$.

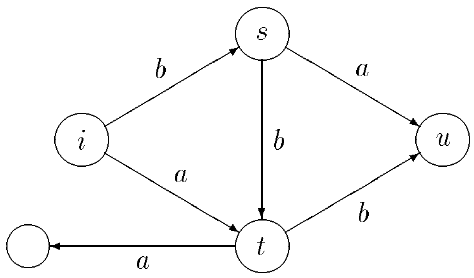

In particular, $\delta(i,x)$ exists and $\delta(i,x) \neq i$ in $\mathcal{B}$. Now it suffices to complete $\mathcal{B}$ to a permutation automaton to produce the required open subgroup K.

2.3.2 General results

It is not true in general that the notions of closed and preclosed coincide.

Let H be a finitely generated subgroup of the free group. Going back to the definition of $\tilde{H}$, it is not difficult to verify that the congruence $\equiv$ on $\mathcal{A}(H)$ coincides with $\sim_{H,Cl(H)}$, using Proposition 2.5. That is, $\mathcal{A}(\tilde{H})$ can be extended to form $\mathcal{A}(Cl(H))$. Therefore $\tilde{H}$ is a free factor of $Cl(H)$ and we have the following result.

Proposition 2.8 *Let H be a finitely generated subgroup of the free group. If H is closed, then H is preclosed.*

We have seen so far that in general, a finitely generated subgroup of the free group is preclosed if it is closed, and it is a free factor of an open subgroup if it is preclosed. Neither implication is an equivalence: in Example 2.3 $\mathbf{V} = \mathbf{G}_{nil}$ and the subgroup bHb^{-1} is a free factor of an open subgroup but it is not preclosed. On the other hand, H is preclosed and, as stated in Example 2.7, it is not closed (Example 3.7).

2.3.3 The extension-closed case

If $\mathbf{V}$ is extension-closed, the notions of closed and preclosed coincide. The key result to prove this fact is Proposition 2.6.

Proposition 2.9 *Let us assume that $\mathbf{V}$ is extension-closed and let H be a finitely generated subgroup of the free group. Then the following are equivalent.*

- *H is closed;*
- *H is preclosed;*

• *H is a free factor of an open subgroup.*

In addition, we have $\tilde{H} = Cl(H)$ and $\mathrm{rank}(Cl(H)) \leq \mathrm{rank}(H)$*.*

The equivalence between the properties of being closed and being a free factor of an open subgroup and the rank inequality were first proved by Ribes and Zalesskiĭ [6].

Remark. It follows from Proposition 2.9 that if $\mathbf{V}$ is extension-closed, then $Cl(H)$ is an overgroup of H. Thus, $Cl(H)$ may take only finitely many values as $\mathbf{V}$ runs over the extension-closed pseudovarieties.

The rank inequality in Proposition 2.9 raises an interesting question. We will see (Section 3.2) that if $\mathbf{V} = \mathbf{G}_{nil}$, then the closure of a finitely generated subgroup H always has finite rank, but that rank may be greater than the rank of H (Example 3.7 below). On the other hand, if $\mathbf{V} = \mathbf{Ab}$, the pseudovariety of finite abelian groups, then the closure of the trivial subgroup is the derived subgroup of $F(A)$, which has infinite rank. However, we do not know any example of a pseudovariety $\mathbf{V}$ giving rise to a Hausdorff topology (that is, the free group is residually $\mathbf{V}$) such that the closure of a finitely generated subgroup might have infinite rank.

3 Practical computations

We saw in Section 2.3.1 that if $\mathbf{V} = \mathbf{G}$, then every finitely generated subgroup is closed. There are two other cases where we know how to effectively compute the closure of a finitely generated subgroup: when $\mathbf{V}$ is the class of all finite p-groups for some prime p, and when $\mathbf{V}$ is the class of all finite nilpotent groups. An important open question, which we briefly discuss at the end of the paper, concerns the case where $\mathbf{V}$ is the class of all finite solvable groups.

3.1 Computing the pro-p closure of a subgroup

By pro-p, we mean pro-$\mathbf{G}_p$, where p is prime and $\mathbf{G}_p$ is the pseudovariety of finite p-groups. We write $Cl_p(H)$ for the pro-p closure of H. The algorithm to compute $Cl_p(H)$ described here, can be found also in [5], with all the necessary justifications. It is adapted from a first algorithm due to Ribes and Zalesskiĭ [6].

The main property of p-groups used here is the well-known fact that in a finite p-group, every maximal proper subgroup has index exactly p [7]. This is used to prove the following lemma.

Lemma 3.1 *Let H be a p-open subgroup of $F(A)$. Then there exists a morphism $\varphi: F(A) \longrightarrow \mathbb{Z}/p\mathbb{Z}$ such that $H \subseteq \ker\varphi$.*

This in turn is used to characterize the p-dense subgroups of $F(A)$, that is, the subgroups which are not contained in a p-open proper subgroup. This requires the consideration of all the morphisms from $F(A)$ into $\mathbb{Z}/p\mathbb{Z}$ (there are only a finite number of them), or equivalently the consideration of the natural morphism σ from $F(A)$ into the free A-generated element of the pseudovariety generated by $\mathbb{Z}/p\mathbb{Z}$, namely $\mathbb{F}_p(A) = (\mathbb{Z}/p\mathbb{Z})^A$. More precisely, we have the following result.

Corollary 3.2 *Let H be a subgroup of $F(A)$. Then H is p-dense if and only if $\sigma(H) = \mathbb{F}_p(A)$.*

Now if we have a basis (or a generating set) of H, then we have one for $\sigma(H)$ too. Since $\sigma(H)$ is also a subspace of the finite-dimensional vector space $\mathbb{F}_p(A)$, deciding whether H is p-dense becomes a simple problem of linear algebra. See Example 3.3 below.

Corollary 3.2 is the basis of the algorithm. In the algorithm, we compute a sequence $H_0 = F(A)$, H_1, $H_2, \ldots$, of p-closed overgroups of H, ending with $Cl_p(H)$.

If H is p-dense in $H_0 = F(A)$, then $Cl_p(H) = F(A)$. Otherwise, $K = \sigma^{-1}(\sigma(H))$ is a proper subgroup of $F(A)$ and a p-open subgroup containing H. Now, if H_1 is given by $\mathcal{A}(H_1) = \mathcal{A}(H)/\sim_{H,K}$, then H_1 is a p-closed overgroup of H and a free factor of K, so $H_1 \neq F(A)$. In addition, it follows from Proposition 2.6 that $Cl_p(H)$ coincides with the pro-p closure of H in the free group H_1. So we compute a basis A_1 of H_1, we express a known basis of H into this new alphabet, and we use again the characterization of p-denseness: either H is p-dense in H_1, $Cl_p(H) = H_1$ and the algorithm stops; or we can compute an overgroup H_2 of H which is p-closed and properly contained in H_1. As long as H is not p-dense in H_k, we continue iterating the process. See Example 3.3 below.

By definition, each H_k constructed in this fashion is such that $\mathcal{A}(H_k) = \mathcal{A}(H)/{\equiv_k}$ for some congruence $\equiv_k$ on $\mathcal{A}(H)$. One shows that $\equiv_k$ is a proper refinement of $\equiv_{k-1}$ for each $k \geq 1$, so the iterative process must stop as there are only finitely many congruences on $\mathcal{A}(H)$. As a result, the last overgroup of H constructed in the process is $Cl_p(H)$.

Note. A precise analysis shows that the computation of $Cl_p(H)$ is performed in time at most $O(n^5)$, where n is the number of positive edges in

$\mathcal{A}(H)$. Since $Cl_p(H)$ is p-preclosed, $\mathcal{A}(Cl_p(H))$ can be extended to a permutation automaton with associated group in $\mathbf{G}_p$: the algorithm described above can be modified to compute such an extension of $\mathcal{A}(Cl_p(H))$, see [5].

Example 3.3 Let $p = 3$ and let H be as in Example 1.1.

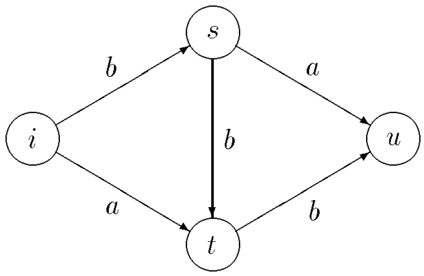

$H = \langle aba^{-1}b^{-1}, b^2a^{-1}\rangle$

The vector space $\sigma(H)$ is generated by the vectors $\sigma(aba^{-1}b^{-1}) = 0 = (0,0)$ and $\sigma(b^2a^{-1}) = 2b - a = (2,2)$, so $\sigma(H) = \langle 0, a+b\rangle \neq \mathbb{F}_3(A)$ and H is not 3-dense in $F(a,b)$.

Now let $K_1 = \sigma^{-1}(\sigma(H))$. We do not need to compute $\mathcal{A}(K_1)$ nor a basis of K_1 in order to compute the congruence $\sim_{H,K_1}$ on $\mathcal{A}(H)$. First we fix a set of geodesics for the states of $\mathcal{A}(H)$, that is, for each state q of $\mathcal{A}(H)$, a reduced word g_q such that $\delta(i, g_q) = q$. In our example, we may set $g_i = 1$, $g_s = b$, $g_t = a$ and $g_u = ba$. Next we observe that if p and q are states, then $p \sim_{H,K_1} q$ if and only if $g_p g_q^{-1} \in K_1$, that is, if and only if $\sigma(g_p g_q^{-1}) \in \sigma(H)$.

Thus we have $i \sim_{H,K_1} u$ and s and t are alone in their respective classes. So $H_1 = \mathcal{H}(\mathcal{A}(H)/{\sim_{H,K_1}})$ is the following group.

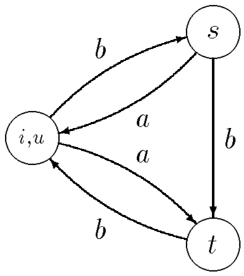

$H_1 = \langle ba, ab, b^3\rangle$

Let us rename the elements of the basis of H_1, say $x = ba$, $y = ab$ and $z = b^3$. It is easy to see that $aba^{-1}b^{-1} = yx^{-1}$ and $b^2a^{-1} = zy^{-1}$, so $H = \langle yx^{-1}, zy^{-1}\rangle$. If σ_1 denotes the natural morphism from $F(x,y,z)$ onto $\mathbb{F}_3(x,y,z)$, we have $\sigma_1(H) = \langle y - x, z - x\rangle$. For dimension reason, it is immediate that $\sigma(H)_1 \neq H_1 = F(x,y,z)$, so H is not 3-dense in H_1.

Let $K_2 = \sigma_1^{-1}(\sigma_1(H))$. We want to compute $\sim_{H,K_2}$, a congruence on $\mathcal{A}(H)$ which refines $\sim_{H,K_1}$. Since $i \sim_{H,K_1} u$ is the only non-trivial relation in $\sim_{H,K_1}$, we consider $\sigma_1(g_i g_u^{-1}) = \sigma_1(ba) = \sigma_1(x) = x = (1,0,0)$. It does not lie in $\sigma_1(H)$, so the congruence $\sim_{H,K_2}$ is the identity relation and $H_2 = H$. Since H is trivially 3-dense in itself, $Cl_3(H) = H$ and H is 3-closed.

This example was given with much detail, even though it was possible to use shortcuts here and there. For instance, we can observe that H_1 has rank greater than H, so H_1 could not be the 3-closure of H. Observing next that there is no non-trivial automaton congruence of $\mathcal{A}(H)$ refining $\sim_{H,K_1}$ leads to the conclusion that H is 3-closed.

Next, an observation made in Example 2.2 shows that H is 2-preclosed and hence 2-closed. In fact, another simple observation shows that H is p-closed for any p. Indeed, the only quotients of $\mathcal{A}(H)$ are itself, the one-state automaton (with two loops), the automaton of the 3-closure of H given above, and a 2-state automaton, where the states are $\{i,t\}$ and $\{s,u\}$. It is easily verified that the two non-trivial quotients give rise to subgroups of $F(A)$ with rank 3, so neither can be a p-closure of H. Thus, for each p, H is either p-closed or p-dense. But one of the generators of H always has a zero image, so we never have $\sigma(H) = \mathbb{F}_p(A)$ and hence, H is never p-dense in $F(A)$.

These observations can be somewhat systematised. Let $h_1, \ldots, h_r$ be a basis of H and let $\mathfrak{M}(H)$ be the $r \times |A|$ matrix in which the j-th row is the commutative image of h_j. This is a matrix with entries in $\mathbb{Z}$. We write $\mathfrak{M}_p(H)$ when the entries are taken in the p-element field.

• If in the above algorithm, we have constructed an overgroup H_k of H with rank greater than $\operatorname{rank}(H)$, then we already know that H cannot be p-dense in H_k and the algorithm will continue for at least one more step.

• If in addition every overgroup of H located properly between H and H_k (that is, every overgroup $\mathcal{H}(\mathcal{A}(H)/\sim)$ where $\sim$ refines the congruence which gave rise to H_k) also has rank greater than $\operatorname{rank}(H)$, then H is p-closed.

• If every overgroup of H properly contained in $F(A)$ has rank greater than $\operatorname{rank}(H)$, then H is p-dense in $F(A)$ if $\mathfrak{M}_p(H)$ has rank $|A|$, and it is p-closed otherwise. This assumption on H was verified in Example 3.3. It is also satisfied, for instance, if $\mathcal{A}(H)$ has no non-trivial congruence: we can call such groups H *primitive* subgroups of $F(A)$ and it would be interesting to have some information about them.

3.2 Computing the pro-nil closure of a subgroup

By pro-nil, we mean pro-$\mathbf{G}_{nil}$, where $\mathbf{G}_{nil}$ is the pseudovariety of finite nilpotent groups. We write $Cl_{nil}(H)$ for the pro-nil closure of H.

The main property of finite nilpotent groups used here is the well-known fact that any finite nilpotent group is a direct product of p-groups [7]. This is used to prove the following result.

Proposition 3.4 *Let H be a finitely generated subgroup of the free group. Then $Cl_{nil}(H) = \bigcap_{p \text{ prime}} Cl_p(H)$. Moreover, $Cl_{nil}(H)$ has finite rank.*

The second part of the statement follows from an earlier remark: since the p-closures of H are overgroups of H, $Cl_{nil}(H)$ is in fact a finite intersection of overgroups of H, each of which is finitely generated, and hence is finitely generated itself. However, the rank of $Cl_{nil}(H)$ may be greater than the rank of H, see Example 3.7.

It follows clearly from Proposition 3.4 that if H is p-closed for some p, then it is nil-closed. Example 3.6 shows that there are nil-closed subgroups which are not p-closed for any p.

An immediate consequence of Proposition 3.4 is that in order to compute $Cl_{nil}(H)$ it suffices to compute all the $Cl_p(H)$: they will take a finite number of values, and we only have to compute their intersection (see Section 1.3.1).

As it turns out, it is possible to simultaneously compute all the p-closures of H, and hence to compute its nil-closure. Having found the nil-closure of H, we can also compute its nil-preclosure: it is the image of $\mathcal{A}(H)$ in $\mathcal{A}(Cl_{nil}(H))$, and it can be obtained by considering the intersection of the congruences of $\mathcal{A}(H)$ which give the automata $\mathcal{A}(Cl_p(H))$. Thus we have the following result.

Proposition 3.5 *It is possible to effectively compute the nil-closure and the nil-preclosure of a finitely generated subgroup of the free group. In particular it is possible to decide whether such a subgroup is nil-closed or nil-preclosed.*

In the rest of this section, we discuss means to compute all the p-closures of H. The case where the matrix $\mathfrak{M}(H)$ has rank $|A|$ is interesting, and it is rapidly reduced to a finite number of computations of p-closures.

In that case, let d be the gcd of the non-zero order $|A|$ minor determinants of $\mathfrak{M}(H)$. For every prime p which does not divide d, H is p-dense and $Cl_p(H) = F(A)$. So we need only to compute the p-closures of H when p divides d, and to take their intersection, see Example 3.6. In the special case where $d = \pm 1$, $Cl_{nil}(H) = F(A)$.

Example 3.6 Let H be the following group.

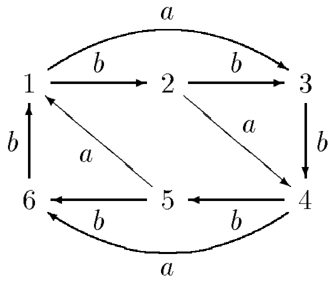

$$H = \langle b^6, ab^4, bab^3, b^3ab, b^4a \rangle$$

$$\mathfrak{M}(H) = \begin{pmatrix} 0 & 6 \\ 1 & 4 \\ 1 & 4 \\ 1 & 4 \end{pmatrix}$$

The gcd of the order 2 non-zero determinants of $\mathfrak{M}(H)$ is 6. As a result, H is p-dense for $p \geq 5$. Running the algorithm of Section 3.1 for $p = 2$ and $p = 3$ shows that the 2- and the 3-closures of H are as follows.

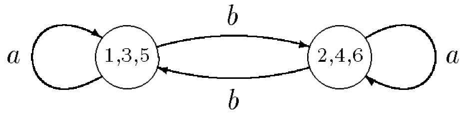

$Cl_2(H) = \langle a, b^2, bab \rangle$

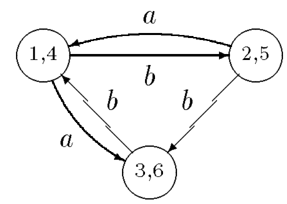

$Cl_3(H) = \langle ba, ab, b^3 \rangle$

In particular H is not p-closed for any p. Now it suffices to take the direct product of these two automata to verify that $H = Cl_2(H) \cap Cl_3(H)$ and hence H is nil-closed.

Example 3.7 Let H and K be as in Example 2.3,

$H = \langle ab^{-1}, a^2b^{-1}a^{-1}, a^3b^{-1}a^{-2}, a^4b^{-1}a^{-3}, a^5b^{-1}a^{-4}, a^6 \rangle$,

$K = \langle ab^{-1}, a^2b^{-1}a^{-1}, a^3b^{-1}a^{-2}, a^4b^{-1}a^{-3}, a^5b^{-1}a^{-4}, a^6, a^5b \rangle$.

We already observed that K is nil-open and H is nil-preclosed. The distinct rows of $\mathfrak{M}(H)$ are $(1, -1)$ and $(6, 0)$. Thus H is p-dense if $p \geq 5$. For $p = 2, 3$, a computation similar to that of Example 3.6 leads to the following result.

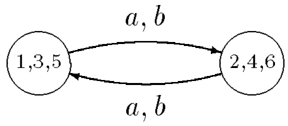

$Cl_2(H) = \langle ab^{-1}, a^2, ab \rangle$

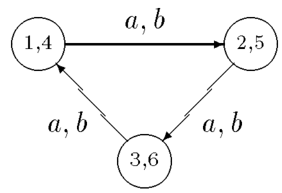

$Cl_3(H) = \langle ba^{-1}, aba^{-2}, a^2b, a^3 \rangle$

The intersection of these two subgroups is easily verified to be exactly K, so $Cl_{nil}(H) = K$ and H is not nil-closed, even though it is nil-preclosed. Note also that $\mathrm{rank}(K) > \mathrm{rank}(H)$.

In addition, $Cl_{nil}(bHb^{-1}) = bKb^{-1} = K$. Since $\mathcal{A}(bHb^{-1})$ cannot be extended to form the automaton associated with $Cl_{nil}(bHb^{-1})$, the subgroup bHb^{-1} is not nil-preclosed, even though it is a free factor of K and it is a conjugate of a nil-preclosed subgroup.

This approach tells us nothing however if $\mathfrak{M}(H)$ has rank less than H. Our only *a priori* information in that case is that H is not p-dense for any p. For instance, for $H = \langle aba^{-1}b^{-1}, b^2a^{-1}\rangle$, $\mathfrak{M}(H)$ has rank 1. We saw in Example 3.3 that H is 3-closed, and hence it is nil-closed. We now describe a general algorithm to compute the nil-closure of H.

The consideration of the minor determinants of $\mathfrak{M}(H)$ allows us to effectively determine the set of primes p such that H is p-dense and it follows from our previous discussion that this set is empty or cofinite. More generally, if H is a subgroup of K, we can compute the set of primes p such that H is p-dense in K.

This observation is then used to compute the set $\mathbb{P}(H)$ of primes p such that H is p-closed. This is done inductively, starting with $H = F\mathcal{G}(A)$ which is p-closed for every p. The induction takes place in the following poset: we consider all the finitely generated subgroups of the free group and we let $K \preceq L$ is K is an overgroup of L. Thus $F\mathcal{G}(A)$ is the least element of this poset. In order to compute $\mathbb{P}(H)$, we need only compute the values of $\mathbb{P}(L)$ when $L \prec H$, that is, we need only consider the quotients of $\mathcal{A}(H)$, which are finitely many. It will turn out to be the case that $\mathbb{P}(H)$ is always finite or cofinite.

Let K be an overgroup of H and assume that we already calculated $\mathbb{P}(L)$ for every overgroup $L \neq K$ of K, and that these sets are finite or cofinite. If p is a prime and K has at least 2 minimal p-closed overgroups L and L', then K must be p-closed since its p-closure is an overgroup of K contained in $L \cap L'$. On the other hand, if K has a single minimal p-closed overgroup L, then the p-closure of K is either L or K, depending on whether H is p-dense in L or not. Thus we can compute $\mathbb{P}(K)$ which is again finite or cofinite.

Recall that our problem is to compute the intersection of the p-closures of H. It suffices to compute the intersection of all the overgroups of H which are p-closed for some prime p. Once the $\mathbb{P}(L)$ have been computed, for all overgroups L of H, we simply need to take the intersection of the overgroups

L such that $\mathbb{P}(L) \neq \emptyset$.

Of course, this is a bit sloppy, and we may take a shorter intersection by first singling out the p-closures of H: the atoms of the Boolean algebra generated by the $\mathbb{P}(L)$ form a finite partition of the set of primes, in finite and cofinite classes. For each prime p, there is a minimal overgroup of H which is p-closed, namely $Cl_p(H)$. In addition, if p and q are in the same partition class, then $Cl_p(H) = Cl_q(H)$. So we can collect the p-closures of H, there will be one for each partition class, and compute their intersection to find the nil-closure of H.

3.3 On the pro-solvable case

The pseudovariety $\mathbf{G}_{sol}$ of finite solvable groups is extension-closed, so we know that the pro-solvable closure of a finitely generated subgroup H is an overgroup of H and it has rank at most $\text{rank}(H)$. Thus we know that this closure is one of a finite computable list of subgroups, but we do not know at the moment how to determine which overgroup of H is $Cl_{sol}(H)$. In fact, the problem is wide open and seems rather difficult.

Here are a few situations where we can say something.

• If the cyclically reduced core of $\mathcal{A}(H)$ has at most 4 states, then H is sol-closed, since every group of permutation on 4 elements is solvable.

• If H is nil-preclosed, then it is sol-closed.

• If H has finite index but is not sol-open, then it is not sol-closed and its sol-closure is computable. Indeed the overgroups of H all have finite index, so their automata cannot be extended non-trivially to a permutation automaton. Thus we need to consider the finitely many permutation automata which are quotients of $\mathcal{A}(H)$ and determine which have a solvable transformation group.

On the other hand, we do not know how to decide whether H is sol-dense in $F(A)$. If we were able to decide this, then we would be able to compute the sol-closure of H like we did for the p-closure.

References

[1] J.-C. Birget, S. Margolis, J. Meakin and P. Weil. PSPACE-completeness of certain algorithmic problems on the subgroups of free groups, *Theoretical Computer Science*, to appear.

[2] M. Hall Jr. A topology for free groups and related groups, *Annals of Mathematics* **52** (1950) 127–139.

[3] R.C. Lyndon and P.E. Schupp. *Combinatorial group theory.* Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 89. Springer, Berlin, 1977.

[4] S. Margolis and J. Meakin. Free inverse monoids and graph immersions, *Int. J. Algebra and Comput.* **3** (1993) 79–100.

[5] S. Margolis, M. Sapir and P. Weil. Closed subgroups in pro-p topologies and the extension problem for inverse automata, to appear.

[6] L. Ribes and P.A. Zalesskiĭ. The pro-p topology of a free group and algorithmic problems in semigroups, *Intern. J. Algebra and Computation* **4** (1994) 359–374.

[7] J. Rotman. *An introduction to the theory of groups* (4th edition), Springer, New York, 1995.

[8] J. Stallings. The topology of graphs, *Inventiones Mathematicæ* **71** (1983) 551–565.