Closed subgroups in pro-V topologies and the extension problem for inverse automata[1]

S. Margolis[2] M. Sapir[3] P. Weil[4]

Abstract We relate the problem of computing the closure of a finitely generated subgroup of the free group in the pro-**V** topology, where **V** is a pseudovariety of finite groups, with an extension problem for inverse automata which can be stated as follows: given partial one-to-one maps on a finite set, can they be extended into permutations generating a group in **V**? The two problems are equivalent when **V** is extension-closed. Turning to practical computations, we modify Ribes and Zalesskiĭ's algorithm to compute the pro-p closure of a finitely generated subgroup of the free group in polynomial time, and to effectively compute its pro-nilpotent closure. Finally, we apply our results to a problem in finite monoid theory, the membership problem in pseudovarieties of inverse monoids which are Mal'cev products of semilattices and a pseudovariety of groups.

Résumé Nous établissons un lien entre le problème du calcul de l'adhérence d'un sous-groupe finiment engendré du groupe libre dans la topologie pro-**V**, où **V** est une pseudovariété de groupes finis, et un problème d'extension pour les automates inversifs qui peut être énoncé de la façon suivante : étant données des transformations partielles injectives d'un ensemble fini, peuvent-elles être étendues en des permutations qui engendrent un groupe dans **V** ? Les deux problèmes sont équivalents si **V** est fermée par extensions. Nous intéressant ensuite aux calculs pratiques, nous modifions l'algorithme de Ribes et Zalesskiĭ pour calculer l'adhérence pro-p d'un sous-groupe finiment engendré du groupe libre en temps polynomial et pour calculer effectivement sa clôture pro-nilpotente. Enfin nous appliquons nos résultats à un problème de théorie des monoïdes finis, celui de de l'appartenance dans les pseudovariétés de monoïdes inversifs qui sont des produits de Mal'cev de demi-treillis et d'une pseudovariété de groupes.

---

[1]The first and third authors received support from the French-Israeli CNRS-MOSA project *Algorithmic problems on automata and semigroups*. The first and the second author were supported by the NSF grant DMS-9623284. All three authors were supported by the Center for Communication and Information Sciences, University of Nebraska-Lincoln.

[2]`margolis@macs.biu.ac.il`. Department of Computer Science – Bar Ilan University – 52900 Ramat Gan, Israel.

[3]`msapir@math.vanderbilt.edu`. Department of Mathematics – Vanderbilt University – Nashville, TN 37240-0001, USA.

[4]`Pascal.Weil@labri.u-bordeaux.fr`. LaBRI, CNRS – 351 cours de la Libération – 33405 Talence Cedex, France.

This paper is at the confluence of several streams of research, namely the theory of profinite groups, the theory of finite monoids and the theory of rational languages.

In [8, 9], Hall defined the pro-$\mathbf{V}$ topology on a group, where $\mathbf{V}$ is a pseudovariety of finite groups (i.e. a class of finite groups closed under taking subgroups, quotients and finite direct products). In the pro-$\mathbf{V}$ topology on a group G, a basis of clopen neighborhoods of 1 is given by the normal subgroups N of G such that N has finite index and $G/N \in \mathbf{V}$. In this paper, we will consider topologies on the free group, and we will refer especially to the cases where $\mathbf{V}$ is the class of all finite groups, the class of all finite p-groups for some prime p, and the class of all finite nilpotent groups. We will talk respectively of the profinite topology, the pro-p topology and the pro-nil topology.

Hall showed that every finitely generated subgroup of the free group is closed in the profinite topology and is a free factor of a clopen subgroup [9]. It is easily seen that this result does not extend to all pro-$\mathbf{V}$ topologies on the free group. However it is a difficult question to decide, for a fixed pseudovariety of groups $\mathbf{V}$, whether a given finitely generated subgroup H is closed, or to compute the closure of H. It is not true in general that the closure of a finitely generated subgroup is finitely generated.

In [18], Ribes and Zalesskiĭ showed that the pro-p closure of a finitely generated subgroup of the free group is finitely generated (a fact which holds whenever the class $\mathbf{V}$ is extension-closed) and effectively computable. They also showed that, in the pro-p topology, closed finitely generated subgroups are free factors of clopen subgroups.

In this paper, we return to Ribes and Zalesskiĭ's paper, both for its theoretical results and for the algorithm to compute the pro-p closure of a finitely generated subgroup. First we attempt to clarify which properties depend on the closure of $\mathbf{V}$ under extension and which do not. We are able to distinguish three properties of a finitely generated subgroup H of the free group, which are equivalent when $\mathbf{V}$ is extension-closed: the first property is for H to be closed in the pro-$\mathbf{V}$ topology; the second property is for H to be a free factor of a clopen subgroup (the clopen subgroups are the finite-index subgroups K such that the quotient of the free group by the core of K – the intersection of all the conjugates of K – is an element of $\mathbf{V}$ [9]); the last property is an extension property which is expressed in terms of the finite state automaton canonically associated with H. This automaton is a finite, graphical representation of the immersion of the bouquet of circles associated with H, see Stallings [21] and the details below. At this point, it is only useful to know that this automaton, or labeled graph, is effectively computable from a set of generators of H, that it uniquely characterizes H, and that it is such that each letter of the alphabet (the set of generators of the free group) labels a partial one-to-one map from the set of states into itself. We say that H is $\mathbf{V}$-extendible if this set of partial one-to-one maps can be extended to a set of permutations of a possibly bigger set of states generating a group in $\mathbf{V}$.

In the first part of the paper, we show that every finitely generated $\mathbf{V}$-closed subgroup is $\mathbf{V}$-extendible, every $\mathbf{V}$-extendible subgroup is a free factor of a $\mathbf{V}$-clopen subgroup, and for any finitely generated subgroup H, there exists a least $\mathbf{V}$-extendible subgroup $\tilde{H}$ containing H. The subgroup $\tilde{H}$ always sits between H and its closure. We also show that H is $\mathbf{V}$-closed if and only if all its conjugates are $\mathbf{V}$-extendible. When $\mathbf{V}$ is extension-closed, we verify that H is $\mathbf{V}$-closed if and only if H is $\mathbf{V}$-extendible, if and only if H is a free factor of a $\mathbf{V}$-clopen subgroup. In particular, in that case, $\tilde{H}$ is the closure of H. Ribes and Zalesskiĭ showed that in the extension-closed case, the rank of the closure of H is always less than or equal to the rank of H [18].

In the second part of the paper, we modify Ribes and Zalesskiĭ's algorithm to compute the pro-p closure of a finitely generated subgroup [18]. The basic idea of the algorithm is not deeply transformed, but we are able to show that our algorithm terminates in polynomial time, whereas the termination of theirs was proved by a compactness argument which did not allow immediate evaluation. The main steps of our algorithm are expressed in terms of computing congruences on the finite automaton associated with H, and solving elementary problems in finite-dimensional linear algebra over the p-element field. Since for the pro-p topology, being closed is equivalent to being extendible, the algorithm also decides whether H is p-extendible. In fact, the algorithm can be modified to compute an extension of the automaton of H consisting of permutations that generate a p-group, or equivalently, computing a p-clopen subgroup K of which H is a free factor.

Next we turn to the pro-nil topology. The previous results do not apply directly here as the class of nilpotent groups is not extension-closed. However, every finite nilpotent group is a direct product of p-groups. We use this well-known fact to prove that the nil-closure of a finitely generated subgroup H is finitely generated and we give an algorithm to effectively compute this closure.

There are important connections between this work and results in the theories of finite monoids and rational languages. One connection was made explicit in Ribes and Zalesskiĭ [18]: ideas originating in [13] are put to work to show how the computation of the pro-p closure of a finitely generated subgroup of the free group can be used to compute the closure of a rational language in the pro-p topology of the free monoid. We refer the reader to [18] for the details, pointing out here that our improved algorithm also speeds the computation of such a closure. We also want to point out that the algorithm to compute the closure of a rational language also uses a very deep result on products of closed subgroups of the free group, which was established by Ribes and Zalesskiĭ as well [17].

This latter result was Ribes and Zalesskiĭ's solution to Rhodes' so-called type II conjecture, which was settled independently by Ash as well [3], see [10]. Rhodes' conjecture had been proposed in the seventies in the context of finite monoid theory. Here too we refer the reader to the survey paper [10] for details. It suffices to mention here that the object of this conjecture was an algorithm to compute the $\mathbf{V}$-kernel of a finite monoid, where $\mathbf{V}$ is a pseudovariety of groups as above. The $\mathbf{V}$-kernel of the monoid M is the set of elements $x \in M$ such that x is always related to the identity in a relational morphism from M into a group in $\mathbf{V}$. Determining whether an element x of M lies in its $\mathbf{V}$-kernel reduces to deciding whether the empty word belongs to the pro-$\mathbf{V}$ closure of a certain rational language.

This computation in turn is linked with the membership problem for pseudovarieties constructed using the Mal'cev product. The last section of the paper solves this problem in the case of Mal'cev products of pseudovarieties of inverse monoids of the form $\mathbf{SL} ⓜ \mathbf{V}$ where $\mathbf{SL}$ is the class of finite semilattices (idempotent and commutative monoids) and $\mathbf{V}$ is a pseudovariety of groups. We prove that the membership problem in such a pseudovariety reduces to deciding whether certain finitely generated subgroups of the free group are $\mathbf{V}$-extendible. A very recent result of Steinberg, which uses the results of this paper that were obtained in preprint form, solves – among other important results – the membership problem for pseudovarieties of monoids of the form $\mathbf{J} ⓜ \mathbf{V}$, where $\mathbf{J}$ is the class of $\mathcal{J}$-trivial monoids [22].

1 The pro-V topology on a group

Here we present general results on profinite topologies on groups. For a general reference on the theory of groups, the reader is referred to [19]. For basic results on profinite groups, see [20, 16, 5, 6]. A more general approach, involving profinite monoids, can be found in [1, 2]. Profinite topologies on groups were introduced by M. Hall [9].

1.1 Definitions

A *pseudovariety of groups* is a class of finite groups closed under taking subgroups, homomorphic images and finite direct products. Important examples are $\mathbf{G}$, the pseudovariety of all finite groups; $\mathbf{G}_p$, the pseudovariety of all finite p-groups (where p is a prime number); $\mathbf{G}_{nil}$ and $\mathbf{G}_{sol}$, the pseudovarieties respectively of all finite nilpotent groups and of all finite solvable groups. If V is a finite group, we let (V) be the pseudovariety generated by V: it is the class of all homomorphic images of subgroups of direct powers V^k of V. The trivial pseudovariety is that which consists only of 1-element groups. In the sequel, we assume that all the pseudovarieties we consider are non-trivial.

Let $\mathbf{V}$ be a pseudovariety of groups and let G be an arbitrary group. The *pro-*$\mathbf{V}$ *topology* on G is defined as the initial topology which makes all morphisms from G into elements of $\mathbf{V}$ continuous (all finite semigroups and groups considered here are equipped with the discrete topology). Thus a basis of open sets is given by the $\varphi^{-1}(x)$, for all morphisms $\varphi\colon G \to V$, with $V \in \mathbf{V}$. Equivalently, G is a topological group and the normal subgroups K of G such that $G/K \in \mathbf{V}$, form a basis of neighborhoods of 1.

We say that two elements $x, y \in G$ can be *separated* by a group V if there exists a morphism $\varphi\colon G \to V$ such that $\varphi x \neq \varphi y$. The group G is *residually* $\mathbf{V}$ if each pair of distinct elements of G, can be separated by some group in $\mathbf{V}$. By definition, the pro-$\mathbf{V}$ topology on G is Hausdorff if and only if G is residually $\mathbf{V}$. Observe that if G is residually $\mathbf{V}$, then all subgroups of G are residually $\mathbf{V}$.

The pro-$\mathbf{V}$ topology on G can also be defined by an ultrametric écart, or quasi-metric, as follows. If $x, y \in G$ can be separated by an element of $\mathbf{V}$, then we let

$$\begin{aligned} r(x,y) &= \min\{|V| \mid V \in \mathbf{V},\ V \text{ separates } x \text{ and } y\}, \\ d(x,y) &= 2^{-r(x,y)}. \end{aligned}$$

If x and y cannot be separated by an element of $\mathbf{V}$, then we let $d(x,y) = 0$. One can verify that the mapping d satisfies

$$d(x,y) = d(y,x) \quad \text{and} \quad d(x,y) \leq \max\left(d(x,z), d(y,z)\right)$$

for all $x, y, z \in G$. For this écart, multiplication in G is contracting, and hence uniformly continuous. Moreover, the topology defined by d on G is exactly the pro-$\mathbf{V}$ topology introduced above.

Let $\hat{G}$ be the Hausdorff completion of (G, d) [4, TG.II.3.7–9]. We say that $\hat{G}$ is the *pro-$\mathbf{V}$ completion* of G. The construction of $\hat{G}$ is as follows. For all $u, v \in G$, we let $u \sim v$ if $d(u, v) = 0$. Then $\sim$ is a congruence on G, and d naturally induces a distance function on the quotient $G/\!\sim$. That is, $(G/\!\sim, d)$ is a metric space, and $G/\!\sim$ is a group whose multiplication is uniformly continuous. By definition, $\hat{G}$ is the completion of that metric space. Of course, if G is residually $\mathbf{V}$, that is, if the écart d on G is in fact a distance function, then $\sim$ is trivial and $\hat{G}$ is the completion of (G, d) in the usual sense. It is important to stress that d, $\sim$ and $\hat{G}$ strongly depend on $\mathbf{V}$, and ought to be written $d_{\mathbf{V}}$, $\sim_{\mathbf{V}}$ and $\hat{G}_{\mathbf{V}}$ respectively.

Let $\imath_G: G \to \hat{G}$ ($\imath$ if G is understood) map each element $x \in G$ to its $\sim$-class. Then $\imath$ is a morphism, onto a dense subset of $\hat{G}$. It is one-to-one if and only if G is residually $\mathbf{V}$, in which case it is usually omitted, that is, we consider G as a dense subset of $\hat{G}$.

A group is said to be *profinite* if it is a projective limit of finite groups. It is known that a group is profinite if and only if it is compact and totally disconnected (see [16]). More generally, we say that a group is pro-$\mathbf{V}$ if it is a projective limit of groups in $\mathbf{V}$. Equivalently, a group is pro-$\mathbf{V}$ if it is profinite and all its finite continuous homomorphic images are in $\mathbf{V}$ (see [16]). Note that a finite group is pro-$\mathbf{V}$ if and only if it belongs to $\mathbf{V}$.

One can show that $\hat{G}$ is (topologically isomorphic to) the projective limit of the homomorphic images of G belonging to $\mathbf{V}$. In other words, the topology of $\hat{G}$ (as the pro-$\mathbf{V}$ completion of G) coincides with the topology it receives as a projective limit. In particular, $\hat{G}$ is pro-$\mathbf{V}$, and hence compact.

If X is a subset of $\hat{G}$, we denote by $\bar{X}$ the topological closure of X in $\hat{G}$. If X is a subset of G, we denote by $Cl(X)$ its topological closure in the pro-$\mathbf{V}$ topology of G. In particular, we have, $Cl(X) = \imath^{-1}(\overline{\imath(X)})$, and $Cl(X) = \bar{X} \cap G$ if G is residually $\mathbf{V}$.

1.2 Properties of morphisms

The following simple observation will be useful in the sequel.

Proposition 1.1 *Let G and G' be groups equipped with their pro-$\mathbf{V}$ topologies. Let $\varphi: G \to G'$ be a group morphism. Then φ is contracting, uniformly* *continuous, and it induces a unique continuous morphism* $\hat{\varphi}: \hat{G} \to \hat{G}'$ *such that* $\hat{\varphi} \circ \imath_G = \imath_{G'} \circ \varphi$.

Proof. Let $u, v \in G$. Every group $V \in \mathbf{V}$ that separates $\varphi(u)$ and $\varphi(v)$, also separates u and v. Thus, by definition of the écarts which define the uniform structures on G and G', the morphism φ satisfies $d(\varphi(u), \varphi(v)) \leq d(u, v)$, and hence φ is uniformly continuous. The rest of the statement follows from the general properties of complete spaces. □

1.3 Open and closed subgroups

We summarize here well-known results on pro-$\mathbf{V}$ topologies (see [9], [6] or [18]). If G is a group and if H is a subgroup of G, the *core* of H is defined to be the greatest normal subgroup H_G of G contained in H. That is, $H_G = \bigcap_{g \in G} g^{-1}Hg$. We let μ_H be the canonical morphism from G onto G/H_G.

Proposition 1.2 [9, Thm 3.1] *Let G be a group equipped with its pro-$\mathbf{V}$ topology, and let H be a subgroup of G. The following are equivalent.*

(1) *H is open;*

(2) *H is clopen;*

(3) *H has finite index and $G/H_G \in \mathbf{V}$.*

Proof. H_G is a normal subgroup of G contained in H, so H is trivially a union of H_G-cosets. In particular, $H = \mu_H^{-1}\mu_H(H)$. So, if $G/H_G \in \mathbf{V}$, then H is clopen.

Observe that H_G has finite index if and only if H has finite index: in one direction it follows from the containment $H_G \subseteq H$ and in the other direction, from the fact that H_G is an intersection of conjugates of H.

If H is open, then H contains an open neighborhood of 1, i.e. a normal subgroup N such that $G/N \in \mathbf{V}$. In particular, $N \subseteq H_G$, so H and H_G have finite index. Moreover, G/H_G is a quotient of G/N, and hence $G/H_G \in \mathbf{V}$. □

Proposition 1.3 [9, Thm 3.3] *Let G be a group equipped with its pro-$\mathbf{V}$ topology, and let H be a subgroup of G. Then $Cl(H)$ coincides with the* *following intersection of open subgroups*

$$Cl(H) = \bigcap_{\substack{K \text{ open subgroup} \\ H \subseteq K}} K = \bigcap_{\substack{\varphi: G \to V \\ V \in \mathbf{V}}} \varphi^{-1}\varphi(H).$$

Proof. $Cl(H)$ is contained in each open subgroup containing H because open subgroups are closed (Proposition 1.2). The first intersection is contained in the second one because each $\varphi^{-1}\varphi(H)$ is open, by definition of the topology.

There remains to see that the second intersection is contained in $Cl(H)$. Let $g \notin Cl(H)$. Then g admits an open neighborhood which avoids H. That is, there exists a morphism $\varphi: G \to V$ such that $V \in \mathbf{V}$ and $\varphi^{-1}\varphi(g) \subseteq G \backslash H$. Then $\varphi(g) \notin \varphi(H)$, so $g \notin \varphi^{-1}\varphi(H)$, and this concludes the proof. $\square$

This yields immediately the following corollary.

Corollary 1.4 *Let G be a group equipped with its pro-$\mathbf{V}$ topology, and let H be a subgroup of G. The following are equivalent.*

(1) *H is closed;*

(2) *H is an intersection of open subgroups;*

(3) *H is the intersection of the open subgroups containing it;*

(4) *H is the intersection of the $\varphi^{-1}(\varphi(H))$, for all morphisms $\varphi: G \to V$ such that $V \in \mathbf{V}$.*

1.4 Pro-V topology of subgroups

If H is a subgroup of the group G, we can consider two topologies on H: its own pro-$\mathbf{V}$ topology as a group, and the restriction to H of the pro-$\mathbf{V}$ topology of G.

Proposition 1.5 *Let G be a group equipped with its pro-$\mathbf{V}$ topology. Let H be a subgroup of G. The restriction to H of the pro-$\mathbf{V}$ topology on G is contained in the pro-$\mathbf{V}$ topology on H.*

Proof. Let U be an elementary open subset of G: there exists a morphism $\varphi: G \to V$ such that $V \in \mathbf{V}$ and $U = \varphi^{-1}(v)$ for some $v \in V$. If ψ is the restriction of φ to H, $\psi = \varphi_{|H}$, then $U \cap H = \psi^{-1}(v)$, so $U \cap H$ is open in the pro-$\mathbf{V}$ topology of H. $\square$

Let H be a subgroup of G and let $\jmath: H \to G$ be the inclusion morphism. We say that H is a *retract* if there exists a morphism $\varphi: G \to H$ such that $\varphi \circ \jmath = \mathrm{id}_H$, i.e. the restriction of φ to H is the identity of H. Equivalently, there exists an onto morphism $\varphi: G \to H$ such that $\varphi^2 = \varphi$. The following situation will arise frequently in the sequel: a free factor of G is a retract of G. Recall that H is a free factor of G if there exists a subgroup H' of G such that G is (isomorphic to) the free product $H * H'$.

The three following statements generalize [18, Lemma 3.1.(ii)].

Proposition 1.6 *Let G be a group equipped with its pro-$\mathbf{V}$ topology. Let H be a retract of G. Then the pro-$\mathbf{V}$ topology of H coincides with the restriction to H of the pro-$\mathbf{V}$ topology of G.*

Proof. Let $\jmath: H \to G$ be the inclusion morphism and let $\varphi: G \to H$ be a morphism such that $\varphi \circ \jmath = \mathrm{id}_H$. By Proposition 1.5, if U is open in the pro-$\mathbf{V}$ topology on G, then $U \cap H$ is open in the pro-$\mathbf{V}$ topology on H.

Conversely, let U be an elementary open set of the pro-$\mathbf{V}$ topology of H. Then there exists a morphism $\psi: H \to V$ such that $V \in \mathbf{V}$ and $U = \psi^{-1}(v)$ for some $v \in V$. Now consider the composite morphism $\psi \circ \varphi: G \to V$. Then $(\psi \circ \varphi)^{-1}(v)$ is open in G, and $(\psi \circ \varphi)^{-1}(v) \cap H = \varphi^{-1}(U) \cap H = U$. So U is open in the restriction to H of the pro-$\mathbf{V}$ topology of G. □

Corollary 1.7 *Let G be a group, equipped with its pro-$\mathbf{V}$ topology, and let H be a retract of G. Let $\imath_G$ be the natural morphism from G into its pro-$\mathbf{V}$ completion $\hat{G}$. Then $\imath_G^{-1}\imath_G(H)$ is closed in G.*

Proof. Let $x \in Cl(\imath_G^{-1}\imath_G(H))$. Then for each $\varepsilon > 0$, there exists an element $h \in \imath_G^{-1}\imath_G(H)$ such that $d(x,h) < \varepsilon$. In particular, $\imath_G(h) \in \imath_G(H)$, so there exists $h' \in H$ such that $\imath_G(h) = \imath_G(h')$, that is, $d(h,h') = 0$. It follows that $d(x,h') < \varepsilon$. Let $\varphi: G \to H$ be an onto morphism such that $\varphi^2 = \varphi$. By Proposition 1.1, we have $d(\varphi(x), \varphi(h')) < \varepsilon$. Now $h' \in H$, so $\varphi(h') = h'$, and thus $d(\varphi(x), h') < \varepsilon$. Since d is ultrametric, we have $d(\varphi(x), x) < \varepsilon$ for each $\varepsilon > 0$, and hence $d(\varphi(x), x) = 0$, that is, $\imath_G(x) = \imath_G(\varphi(x))$. But $\varphi(x) \in H$, so $x \in \imath_G^{-1}\imath_G(H)$. □

In the case of a residually $\mathbf{V}$ group, Corollary 1.7 yields the following.

Corollary 1.8 *Let G be a residually $\mathbf{V}$ group, equipped with its pro-$\mathbf{V}$ topology, and let H be a retract of G. Then H is closed in G.*

Proof. Under the assumption that G is residually $\mathbf{V}$, the morphism $\imath_G$ is one-to-one. By convention we can omit it, that is, we may consider G as a subgroup of $\hat{G}$. The result then follows immediately from Corollary 1.7. $\square$

We say that *the finite quotients of G are closed under $\mathbf{V}$-extension* if, whenever K is a finite quotient of G, N is a normal subgroup of K, $N \in \mathbf{V}$ and $K/N \in \mathbf{V}$, then $K \in \mathbf{V}$. Naturally, this is the case if the class $\mathbf{V}$ is closed under extension. The following statement is a restatement of [7, Lemma 3.1].

Proposition 1.9 *Let G be a group equipped with the pro-$\mathbf{V}$ topology. If the finite quotients of G are closed under $\mathbf{V}$-extension, and if H is a clopen subgroup of G, then the pro-$\mathbf{V}$ topology of H coincides with the restriction to H of the pro-$\mathbf{V}$ topology of G.*

Proof. By Proposition 1.5, we know that any open subset of G which is contained in H, is also open in the pro-$\mathbf{V}$ topology on H.

Conversely, let us assume that $U \subseteq H$ is an elementary open set of the pro-$\mathbf{V}$ topology of H. There exists a morphism $\varphi\colon H \to V$ such that $V \in \mathbf{V}$ and $U = \varphi^{-1}(v)$ for some $v \in V$. Up to a translation, we may assume that $U = \varphi^{-1}(1)$, that is, U is a normal subgroup of H and $V = H/U$. In particular, U has finite index in H. By Proposition 1.2, H has finite index in G, so U has finite index in G. By Proposition 1.2, we need to prove that $G/U_G \in \mathbf{V}$.

Since H has finite index in G, we have $G = \bigcup_{i=1}^{r} Hk_i$ for some finite family $k_1, \ldots, k_r$ of elements of G. In particular, $H_G = \bigcap_{i=1}^{r} k_i^{-1}Hk_i$. We also have $U_G = \bigcap_{i=1}^{r} k_i^{-1}Uk_i$ since U is normal in H. Now U_G is normal in G and $U_G \subseteq H_G$, so

$$G/H_G = (G/U_G)/(H_G/U_G).$$

We know that $G/H_G \in \mathbf{V}$, so it suffices to prove that $H_G/U_G \in \mathbf{V}$. Let σ be the morphism

$$\begin{aligned} \sigma\colon \bigcap_{i=1}^{r} k_i^{-1}Hk_i \quad &\longrightarrow \quad \prod_{i=1}^{r}(k_i^{-1}Hk_i/k_i^{-1}Uk_i) \\ x \quad &\longmapsto \quad (x(k_i^{-1}Uk_i))_{1\leq i\leq r} \end{aligned}$$

Then $\ker\sigma = U_G$. Moreover, $\prod_{i=1}^{r}(k_i^{-1}Hk_i/k_i^{-1}Uk_i)$ is isomorphic to the direct product of r copies of H/U. Therefore H_G/U_G is isomorphic to a subgroup of $(H/U)^r$, and hence H_G/U_G lies in $\mathbf{V}$. This concludes the proof. $\square$

Example 1.10 In the above proposition, the hypothesis that the finite quotients of G are closed under $\mathbf{V}$-extension is needed. Indeed, suppose that $\mathbf{V}$ is the class of nilpotent groups and G is the free group over the 2-letter alphabet $\{a,b\}$: the finite quotients of G are not closed under $\mathbf{V}$-extension.

Let S_3 be the symmetric group on 3 elements and let $\varphi: G \to S_3$ be the morphism given by $\varphi(a) = (12)$ and $\varphi(b) = (13)$. Let $K = \ker\varphi$. Let $\varepsilon: S_3 \to \mathbb{Z}_2$ be the signature morphism and let $H = \ker(\varepsilon \circ \varphi)$. Then H and K are normal in G and $K \subseteq H$. Moreover, $G/H = \mathbb{Z}_2$, so H is clopen in the pro-nilpotent topology of G, and $H/K = \mathbb{Z}_3$, so K is clopen in the pro-nilpotent topology of H. But $G/K = S_3$ is not nilpotent, so K is not clopen in the pro-nilpotent topology of G. Thus the conclusion of Proposition 1.9 does not hold for the pro-nilpotent topology on G.

We will see (Example 3.8) that the closure of K in G is H.

2 The case of free groups

Let $F(A)$ be the free group over a fixed finite alphabet A. We will denote by $\hat{F}_A(\mathbf{V})$ the pro-$\mathbf{V}$ completion of $F(A)$.

The group $\hat{F}_A(\mathbf{V})$ enjoys certain freeness properties. Firstly, the image $\imath(F(A))$ of $F(A)$ in $\hat{F}_A(\mathbf{V})$ is the free object over A of the variety of groups generated by the (finite) groups in $\mathbf{V}$. For instance, if $\mathbf{V}$ is the pseudovariety of finite abelian groups, then $\imath(F(A))$ is the free abelian group over A. Of course, if the free group is residually $\mathbf{V}$ (e.g. $\mathbf{V} = \mathbf{G}_p$, $\mathbf{G}_{nil}$, $\mathbf{G}_{sol}$), then $\imath$ is injective, and $\imath(F(A))$ is the free group over A.

Moreover, $\hat{F}_A(\mathbf{V})$ is the *free pro-*$\mathbf{V}$ *group over* A (see for instance [2, Proposition 1.3]):

Proposition 2.1 *Let $\sigma: A \to H$ be a mapping into a pro-$\mathbf{V}$ group H. Then there exists a unique continuous morphism $\hat{\sigma}: \hat{F}_A(\mathbf{V}) \to H$ such that $\hat{\sigma} \circ \imath = \sigma$. If $\sigma(A)$ generates a dense subgroup of H, then $\hat{\sigma}$ is onto.*

A simple situation arises when $\mathbf{V}$ has a finite free object $F_A(\mathbf{V})$ over the alphabet A (for instance if $\mathbf{V}$ is generated by a single group).

Proposition 2.2 *Let $\mathbf{V}$ have a finite free object $F_A(\mathbf{V})$. Let σ be the natural projection $\sigma: F(A) \to F_A(\mathbf{V})$. For each subset $X \subseteq F(A)$, we have $Cl(X) = \sigma^{-1}(\sigma(X))$.*

Proof. Every morphism $\psi: F(A) \to V$ into a group $V \in \mathbf{V}$ can be factorized through σ. The result then follows immediately from Proposition 1.3. $\square$

But there is much more to be gained from the freeness of $F(A)$: we now turn to the representation of finitely generated subgroups of $F(A)$ by means of finite automata.

2.1 Inverse automata and subgroups of the free group

For the results not proved in this section, the reader is referred to [21]. The central idea is a graphical representation of ideas that go back to the early part of the twentieth century (see [11] and [19, Chapter 11]).

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 an A-labeled graph $\mathcal{A}(H)$ in three steps.

First, we construct a set of $|Y|$ loops around a common distinguished vertex 1, 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 last operation consists in "reducing" the graph: we iteratively remove vertices of degree one other than 1.

Example 2.3 $Y = \{bab^{-1}, b^2aa^{-2}\}$. Some steps of the computation.

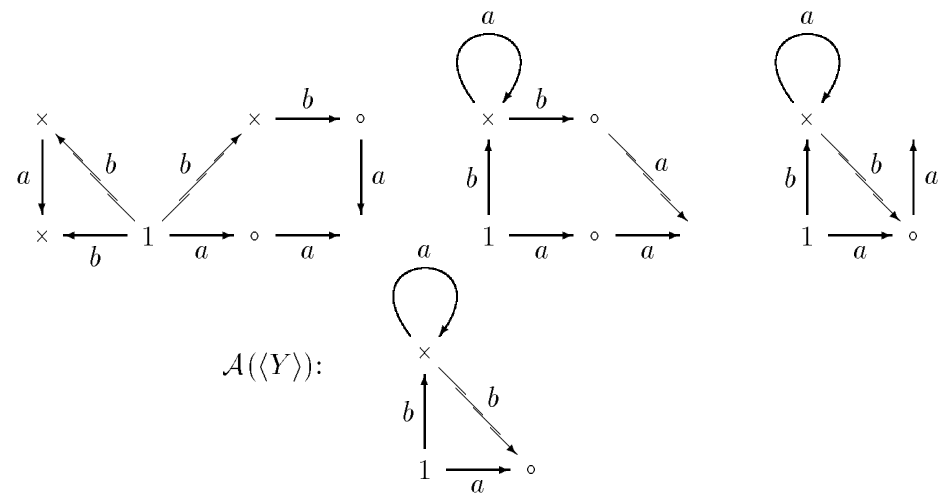

Observe that this construction is performed in time $O(n^2)$, where n is the length of the input, that is, n is the sum of the lengths of the elements of Y.

The labeled graph $\mathcal{A}(H)$ thus constructed is determined by H, and does not depend on the choice of the particular generating set Y [11]; this justifies the notation $\mathcal{A}(H)$.

An *inverse automaton* over an alphabet A is a tuple $\mathcal{A} = (Q, A, \delta, i, f)$ where Q is a finite set, called the set of states, $i \in Q$ is the initial state and $f \in Q$ is the final state. $\delta : Q \times (A \cup A^{-1}) \to Q$ is a partial function such that $\delta(p, a) = q$ if and only if $\delta(q, a^{-1}) = p$ for each $a \in A$ and $p, q \in Q$. The mapping δ is called the transition function of $\mathcal{A}$, and it is usually written $\delta(q, a) = q \cdot a$ $(q \in Q,\ a \in A \cup A^{-1})$. We say that the automaton $\mathcal{A}$ is *complete* if for each $q \in Q$ and for each $a \in A \cup A^{-1}$, $q \cdot a$ exists.

Let $(A \cup A^{-1})^*$ be the free monoid over $A \cup A^{-1}$, that is, the set of finite sequences of letters of $A \cup A^{-1}$. The transition function is extended to $Q \times (A \cup A^{-1})^*$ by letting $q \cdot 1 = q$ and $q \cdot (ua) = (q \cdot u) \cdot a$ (if this is defined) for all $q \in Q$, $u \in (A \cup A^{-1})^*$ and $a \in A \cup A^{-1}$. The *transition morphism* of $\mathcal{A}$ is the morphism defined on $(A \cup A^{-1})^*$ which maps each word u to the partial tranformation of Q given by $q \mapsto q \cdot u$. The *transition monoid* of $\mathcal{A}$ is the range of the transition morphism, that is, the monoid of partial transformations of Q induced by the words of $(A \cup A^{-1})^*$.

We say that a state $q \in Q$ *has degree 1* if

$$\mathrm{Card}\{p \in Q \mid \exists a \in A \cup A^{-1},\ q \cdot a = p\} = 1.$$

We say that $\mathcal{A}$ is *reduced* if its initial state i coincides with its final state f, if every state lies on a path from i to i, and if no state has degree 1, except maybe the initial-final state i.

If H is a finitely generated subgroup of $F(A)$, we can view $\mathcal{A}(H)$ as a reduced inverse automaton over A as follows: the vertices of $\mathcal{A}(H)$ are considered as states, with the distinguished vertex 1 as initial-final state. If $a \in A$ labels an edge from vertex p to q, we let $p \cdot a = q$ and $q \cdot a^{-1} = p$. By construction, $\mathcal{A}(H)$ is a reduced inverse automaton. It is easily verified that if u is a reduced word in $(A \cup A^{-1})^*$, then $1 \cdot u = 1$ if and only if $u \in H$ and $\mathcal{A}(H)$ is the minimal automaton with this property [11].

There is also a reverse construction, from a reduced inverse automaton to a finitely generated subgroup of $F(A)$. More precisely, if $\mathcal{A} = (Q, A, \delta, i, i)$ is a reduced inverse automaton, one can effectively construct a basis of a finitely generated subgroup H of $F(A)$ such that $\mathcal{A} = \mathcal{A}(H)$.

First we compute a spanning tree T of the graph $\mathcal{A}$. For each state q of $\mathcal{A}$, there is a unique shortest path from i to q within T: we let u_q be the label (in $(A \cup A^{-1})^*$) of this path. 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 $y_j = u_{p_j} a_j u_{q_j}^{-1} \in (A \cup A^{-1})^*$, and let $H = \langle y_1, \ldots, y_k \rangle$. Then $\{y_1, \ldots, y_k\}$ is a basis for H and $\mathcal{A} = \mathcal{A}(H)$ [21].

Example 2.4 Let $H = \langle a^2b^{-1}, ab^{-2}, ba, a^{-1}baba^{-1} \rangle$. Then a^3, ba^{-2} and a^2ba^{-1} form a basis of H.

$\mathcal{A}(H)$:

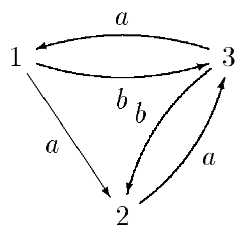

a spanning tree T:

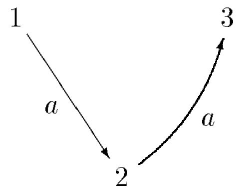

Observe that the construction of a spanning tree can be performed in time $O(e)$, where e is the number of edges of the automaton $\mathcal{A}$. If we want to effectively write down a basis of H, we need to find the geodesics in T from 1 to each vertex, and then to write down one generator for each edge not in the spanning tree. This takes time $O(en)$, where n is the number of vertices of $\mathcal{A}$. In particular, we can combine the two constructions given above to compute, in quadratic time, a basis of the subgroup generated by a given finite set of elements of $F(A)$.

Note. There are analogues of the constructions above, which are valid for any subgroup H of the free group. If we start from an infinite set Y of generators of H, the resulting labeled graph $\mathcal{A}(H)$ again depends only on H, but it may be an infinite graph (it is infinite if and only if H does not have finite rank). That graph can again be considered as an inverse automata, with a possibly infinite set of states.

An *inverse monoid* is a monoid M such that for each element $x \in M$, there exists a unique element $x' \in M$ such that $xx'x = x$ and $x'xx' = x'$. The element x' is called the *inverse* of x, and it is written $x' = x^{-1}$. Inverse monoids form a variety of algebras, and we denote by $FIM(A)$ the free inverse monoid on A. See [12] for basic facts about inverse monoids.

If $\mathcal{A} = (Q, A, i, \delta, f)$ is an inverse automaton, then each letter $a \in A$ induces a partial one-to-one transformation on the set Q, and letter a^{-1} induces the reciprocal transformation. It easily follows that the transition monoid of $\mathcal{A}$ is an inverse monoid. If H is a finitely generated subgroup of $F(A)$, we denote by μ_H and $M(H)$ respectively the transition morphism and the transition monoid of $\mathcal{A}(H)$.

It is well-known that H has finite index if and only if $\mathcal{A}(H)$ is a complete automaton (and hence each letter labels a permutation of the states), if and only if $M(H)$ is a group [21]. In that case, the set of vertices of $\mathcal{A}(H)$ is in bijection with the set of right cosets of H, and $\mathcal{A}(H)$ describes the action of $F(A)$ on these cosets by right translation. In particular, $M(H) = F(A)/H_{F(A)}$. By Proposition 1.2, this implies the following.

Proposition 2.5 *Let H be a finitely generated subgroup of the free group $F(A)$. Then H is clopen if and only if $\mathcal{A}(H)$ is complete and $M(H) \in \mathbf{V}$.*

An *automaton morphism* between inverse automata with one initial-terminal state (from now on always denoted 1) is a mapping between the sets of states, which takes 1 to 1, and which preserves the labeled transitions. An *automaton congruence* on the automaton $\mathcal{A}$ is an equivalence relation $\sim$ on Q, the set of states of $\mathcal{A}$, such that if $p \sim q$ and $p \cdot a$ and $q \cdot a$ exist in $\mathcal{A}$ ($a \in A \cup A^{-1}$), then $p \cdot a \sim q \cdot a$. The quotient automaton $\mathcal{A}/\!\sim$ has set of states $Q/\!\sim$, it has initial-terminal state the $\sim$-class of 1, and it has an a-labeled transition from $[p]$ to $[q]$ if there exists an a-labeled transition of $\mathcal{A}$ from p' to q' for some states $p' \sim p$ and $q' \sim q$. If $\varphi: \mathcal{A} \to \mathcal{B}$ is an automaton morphism, and if $\sim$ is the induced congruence on $\mathcal{A}$ (two states p and q are $\sim$-related if and only if $\varphi(p) = \varphi(q)$), then there is a natural embedding of $\mathcal{A}/\!\sim$ into $\mathcal{B}$, which is onto if φ is.

Proposition 2.6 *Let H and H' be finitely generated subgroups of $F(A)$. Then H is contained in H' if and only if there exists an automaton morphism from $\mathcal{A}(H)$ into $\mathcal{A}(H')$.*

That morphism, if it exists, is unique. Moreover, if it is one-to-one, then H is a free factor of H'.

Proof. Suppose that $H \subseteq H'$: if B is a generating set of H, we can add to B a finite number of elements of H' to get a generating set B' of H'. Now consider the construction process of $\mathcal{A}(H')$ from B': clearly $\mathcal{A}(H')$ can be obtained from $\mathcal{A}(H)$ by adding vertices and transitions, and then identifying states. So there is an automaton morphism from $\mathcal{A}(H)$ into $\mathcal{A}(H')$.

Conversely, let $\varphi: \mathcal{A}(H) \to \mathcal{A}(H')$ be an automaton morphism. If a reduced word u lies in H, then $1 \cdot u = 1$ in $\mathcal{A}(H)$, so, via φ, $1 \cdot u = 1$ in $\mathcal{A}(H')$, and hence $u \in H'$.

The uniqueness is immediately derived from the fact that the automata are inverse and that we impose $\varphi(1) = 1$.

Let us now assume that $\varphi: \mathcal{A}(H) \to \mathcal{A}(H')$ is one-to-one. Any spanning tree of $\mathcal{A}(H)$ (identified with its image under φ) can be extended to a spanning tree of $\mathcal{A}(H')$. It follows from the basis construction procedure given above that there exists a basis of H' which contains a basis for H, and hence H is a free factor of H'. $\square$

Observe that there are many more free factors of H' than can be represented by sub-automata of $\mathcal{A}(H')$. For instance, the free group on two generators is freely generated by ab and b, so $\langle ab \rangle$ is a free factor of the free group, although its automaton does not embed in the bouquet of two circles. Only finitely many free factors of H' arise as sub-automata of $\mathcal{A}(H')$. Of course, the automaton $\mathcal{A}(H')$ is essentially dependent on the choice of A, the set of generators of the free group.

Similarly, there are only finitely many subgroups H' of the free group such that the morphism from $\mathcal{A}(H)$ to $\mathcal{A}(H')$ is onto, or equivalently, such that $\mathcal{A}(H') = \mathcal{A}(H)/\sim$ for some congruence $\sim$ on $\mathcal{A}(H)$. Such subgroups are called *overgroups* of H.

If H is a finitely generated subgroup of $F(A)$ and if $H \subseteq K$, we let $\sim_{H,K}$ be the automaton congruence on $\mathcal{A}(H)$ induced by the morphism from $\mathcal{A}(H)$ into $\mathcal{A}(K)$. Suppose that, for each state p of $\mathcal{A}(H)$, u_p is a reduced word such that $1 \cdot u_p = p$ in $\mathcal{A}(H)$. Then two states p and q of $\mathcal{A}(H)$ are $\sim_{H,K}$-equivalent if and only if $u_p u_q^{-1} \in K$.

2.2 Extendible subgroups

We say that a subgroup H of the free group $F(A)$ is $\mathbf{V}$-*extendible* (extendible, if the pseudovariety $\mathbf{V}$ is understood) if its automaton can be embedded into a complete automaton with transition group in $\mathbf{V}$, that is, into the automaton of a clopen subgroup (Proposition 2.5). In other words, H is extendible if and only if there exists a clopen subgroup K such that $\sim_{H,K}$ is the trivial congruence.

Proposition 2.7 *Let H be a finitely generated subgroup of $F(A)$. If H is extendible, then H is a free factor of a clopen subgroup.*

Proof. This is a trivial consequence of Proposition 2.6. $\square$

We will see (Example 2.15) that the converse is not true, that is, there are free factors of clopen subgroups which are not extendible. However the two notions coincide if the pseudovariety of groups $\mathbf{V}$ is extension-closed (Proposition 2.17).

We will examine in more detail the relationship between the property of extendibility and the topological properties of a subgroup in Sec. 2.3. Let us first remark that each finitely generated subgroup of $F(A)$ is contained in a "best" extendible subgroup. If H is a finitely generated subgroup of the free group, let $\sim$ be the intersection of the $\sim_{H,K}$, where the intersection runs over all clopen subgroups K containing H.

Lemma 2.8 *Let H be a finitely generated subgroup of $F(A)$. Then the automaton congruence $\sim$ coincides with $\sim_{H,Cl(H)}$. Moreover, there exists a clopen subgroup K containing H such that $\sim$ coincides with $\sim_{H,K}$.*

Proof. For each state p of $\mathcal{A}(H)$, let u_p be a reduced word such that $1 \cdot u_p = p$ in $\mathcal{A}(H)$. Let p, q be states of $\mathcal{A}(H)$. Then $p \sim_{H,Cl(H)} q$ if and only if $u_p u_q^{-1} \in Cl(H)$. But $Cl(H)$ is the intersection of the clopen subgroups containing H, so $\sim_{H,Cl(H)}$ and $\sim$ coincide.

Since $\mathcal{A}(H)$ is finite, there are only finitely many congruences on $\mathcal{A}(H)$, so there exists a finite collection of clopen subgroups containing H, say $K_1, \ldots, K_r$, such that $\sim = \bigcap_{i=1}^{r} \sim_{H,K_i}$. It follows from the definition of the $\sim_{H,K_i}$ that $\sim = \sim_{H,K}$, where $K = \bigcap_{i=1}^{r} K_i$. Since K is a finite intersection of clopen subgroups of $F(A)$, K is clopen. □

Let $\tilde{H}$ be the subgroup of $F(A)$ such that $\mathcal{A}(\tilde{H}) = \mathcal{A}(H)/\sim$. Equivalently, $\mathcal{A}(\tilde{H})$ is the image of $\mathcal{A}(H)$ in $\mathcal{A}(Cl(H))$. By definition, $\tilde{H}$ is finitely generated.

Proposition 2.9 *Let H be a finitely generated subgroup of $F(A)$. The subgroup $\tilde{H}$ is the least extendible subgroup containing H.*

Proof. By Lemma 2.8, $\sim = \sim_{H,K}$ for some clopen subgroup K containing H. Therefore $\mathcal{A}(\tilde{H})$ embeds in $\mathcal{A}(K)$, and hence $\tilde{H}$ is extendible.

Conversely, let H' be an extendible subgroup containing H and let K be a clopen subgroup such that $\mathcal{A}(H')$ embeds in $\mathcal{A}(K)$. Then K contains H, so $\sim$ is contained in $\sim_{H,K}$. We want to show that $\sim$ is contained in $\sim_{H,H'}$.

For each state p of $\mathcal{A}(H)$, let u_p be a reduced word such that $1 \cdot u_p = p$ in $\mathcal{A}(H)$. Since H is contained in H', there exists an automaton morphism from $\mathcal{A}(H)$ into $\mathcal{A}(H')$, so u_p also labels a path in $\mathcal{A}(H')$ starting at state 1. Now $\mathcal{A}(H')$ embeds into $\mathcal{A}(K)$, so if p and q are states of $\mathcal{A}(H)$, then

$p \sim_{H,K} q$ if and only if $1 \cdot u_p = 1 \cdot u_q$ in $\mathcal{A}(K)$, if and only if the same equality holds in $\mathcal{A}(H')$, that is, if and only if $p \sim_{H,H'} q$. Therefore $p \sim q$ implies $p \sim_{H,K} q$ implies $p \sim_{H,H'} q$. Thus there exists a morphism from $\mathcal{A}(\tilde{H})$ into $\mathcal{A}(H')$, and hence $\tilde{H} \subseteq H'$. $\square$

This leads to the following properties of extendible subgroups.

Corollary 2.10 *Let H be a finitely generated subgroup of $F(A)$. H is extendible if and only if $\sim$ is trivial on $\mathcal{A}(H)$, if and only if $\tilde{H} = H$.*

Proof. By definition of $\tilde{H}$, $\sim$ is trivial on $\mathcal{A}(H)$ if and only if $\tilde{H} = H$. By Proposition 2.9, H is extendible if and only if $\tilde{H} = H$. $\square$

Corollary 2.11 *Let H be a finitely generated subgroup of $F(A)$. Then $\tilde{\tilde{H}} = \tilde{H}$.*

2.3 Closed subgroups

We first verify that if H is a finitely generated subgroup of $F(A)$, then $\tilde{H}$ sits between H and its closure.

Proposition 2.12 *Let H be a finitely generated subgroup of $F(A)$. Then $H \subseteq \tilde{H} \subseteq Cl(H)$ and $\tilde{H}$ is a free factor of $Cl(H)$.*

Proof. It suffices to observe that, by Lemma 2.8, the congruences $\sim$ and $\sim_{H,Cl(H)}$ on $\mathcal{A}(H)$ coincide, so $\mathcal{A}(\tilde{H}) = \mathcal{A}(H)/\sim$ embeds in $\mathcal{A}(Cl(H))$. $\square$

The second part of the following statement is [18, Corollary 3.1.(i)].

Corollary 2.13 *Every finitely generated closed subgroup is extendible, and is a free factor of a clopen subgroup.*

Note. It is clear that if H is a finitely generated subgroup of $F(A)$, then the subgroup $\tilde{H}$ has finite rank, since its automaton $\mathcal{A}(\tilde{H})$ is a quotient of the finite automaton $\mathcal{A}(H)$. It is not true however that $Cl(H)$ has finite rank in general. $H = \{1\}$ has rank 0 and if $\mathbf{V} = \mathbf{Ab}$, the pseudovariety of finite abelian groups, then $Cl(H)$ is the derived subgroup of $F(A)$, which has infinite rank. See Propositions 2.18, 2.19 and 4.3 below.

Example 2.14 Let $\mathbf{V}$ be the pseudovariety of nilpotent groups. We exhibit an extendible subgroup which is not closed. Let K be the subgroup of $F(A)$ whose automaton is:

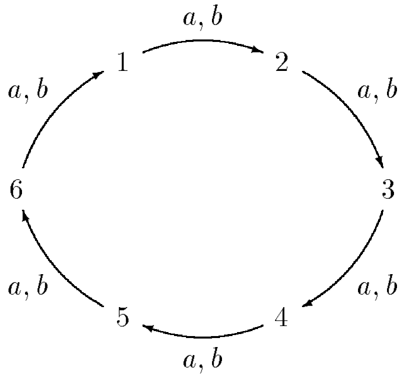

$$K = \langle ab^{-1}, a^2b^{-1}a^{-1}, a^3b^{-1}a^{-2}, \\ a^4b^{-1}a^{-3}, a^5b^{-1}a^{-4}, a^6, a^5b\rangle$$

K is clopen since $M(K)$ is the cyclic 6-element group, which is trivially nilpotent. Now let H be the group whose automaton is obtained from that of K by removing one arrow as follows:

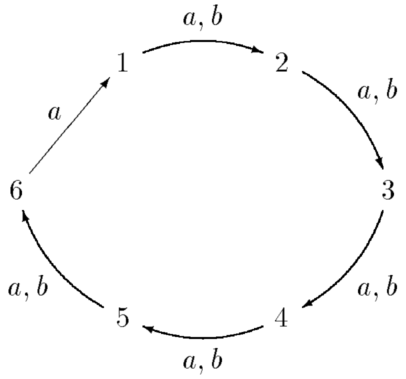

$$H = \langle ab^{-1}, a^2b^{-1}a^{-1}, a^3b^{-1}a^{-2}, \\ a^4b^{-1}a^{-3}, a^5b^{-1}a^{-4}, a^6\rangle$$

H is trivially extendible since its automaton embeds in that of K. But one can show that the closure of H is K (see Example 4.5). Note that H has rank 6 and K has rank 7, that is, the closure of H has rank greater than that of H (see Proposition 2.19).

Example 2.15 A slight modification of the previous example shows that a finitely generated subgroup may be a free factor of a clopen subgroup, yet not be extendible. We use the notation of Example 2.14. Let $H' = bHb^{-1}$.

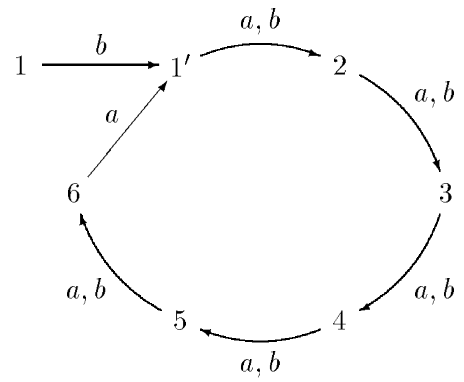

Observe that K is normal, so $bKb^{-1} = K$. Since conjugation by b is a homeomorphism, H' is a free factor of K and $Cl(H') = K$, but $\mathcal{A}(H')$ clearly does not embed in $\mathcal{A}(K)$, that is, H' is not extendible.

Note. This shows that the property of extendibility for a subgroup of the free group is a combinatorial property, but it is neither purely topological nor purely algebraic, since it is not preserved by conjugation.

The next result establishes a weak converse of the second statement of Corollary 2.13. It gives a necessary and sufficient condition for a finitely generated subgroup to be closed, which is inspired by M. Hall's theorem which states that finitely generated subgroups of the free group are closed [8], or at least by its proof in [21].

Proposition 2.16 *Let H be a finitely generated subgroup of $F(A)$. Then H is closed if and only if all the conjugates of H are extendible.*

Proof. In this proof, we let Q be the set of states of $\mathcal{A}(H)$, we let $\cdot$ be the transition function of $\mathcal{A}(H)$, and we let $1 \in Q$ be the designated vertex of $\mathcal{A}(H)$. For short, we write $\mathcal{A}(H) = (Q, 1, \cdot)$.

Conjugation by an element $x \in F(A)$ is a homeomorphism, so if H is closed, then each conjugate $x^{-1}Hx$ is closed, and hence extendible by Corollary 2.13.

To prove the converse, we assume that all conjugates of H are extendible. We first consider the situation where $\mathcal{A}(H)$ does not have any vertex of degree 1 (that is, the vertex 1 of $\mathcal{A}(H)$ does not have degree 1). We will show that for each element $x \in F(A) \setminus H$, there exists a clopen subgroup K containing H, and such that $x \notin K$. This suffices to show that each such x lies outside $Cl(H)$, and hence that $H = Cl(H)$.

Let x be a reduced word in $F(A) \setminus H$. If x labels a path in $\mathcal{A}(H)$ starting at 1, let $q = 1 \cdot x$. Then $1 \neq q$ since $x \notin H$. Since H is extendible, the automaton $\mathcal{A}(H)$ embeds in the automaton $\mathcal{A}(K)$ of a clopen subgroup K. In $\mathcal{A}(K)$, $1 \cdot x = q$, as in $\mathcal{A}(H)$, so $x \notin K$.

Let us now assume that x does not label a path starting at 1 in $\mathcal{A}(H)$. Let x_1 be the longest prefix of x which can be read in $\mathcal{A}(H)$ starting at 1, and let x_2 be the reduced word such that $x = x_1x_2$. We extend $\mathcal{A}(H)$ with $|x_2|$ new edges and vertices to form a path reading x_2 starting at $1 \cdot x_1$. Notice that this extended automaton is again an inverse automaton. We let Q^+ be the extended set of vertices, and we still write $\cdot$ for the transition function of the extended automaton. We let $q = 1 \cdot x = (1 \cdot x_1) \cdot x_2$. By construction, q is not a state of $\mathcal{A}(H)$, and it is the only element of Q^+ which is of degree 1.

We claim that, if H' is the subgroup such that $\mathcal{A}(H') = (Q^+, q, \cdot)$, then $H' = x^{-1}Hx$. Indeed, if u is a reduced word which labels a path from q to q in $\mathcal{A}(H')$, then u must be of the form $x_2^{-1}vx_2$, where v is a reduced word which labels a path from $1 \cdot x_1$ to $1 \cdot x_1$ within $\mathcal{A}(H)$. We have $(1 \cdot x_1) \cdot v = 1 \cdot x_1$, and hence $1 \cdot (x_1vx_1^{-1}) = 1$ in $\mathcal{A}(H)$. So $x_1vx_1^{-1} \in H$, that is, $v \in x_1^{-1}Hx_1$ and $u = x_2^{-1}vx_2 \in x_2^{-1}x_1^{-1}Hx_1x_2 = x^{-1}Hx$. Conversely, let u be a reduced word in H. Then $1 \cdot u = 1$, so $q \cdot x^{-1}ux = (1 \cdot x) \cdot x^{-1}ux = 1 \cdot x = q$. This establishes the claim.

Since $x^{-1}Hx$ is extendible, $\mathcal{A}(H')$ embeds in the automaton $\mathcal{A}(K)$ of some clopen subgroup K. Since $\mathcal{A}(H)$ is contained in $\mathcal{A}(H')$, the automaton $\mathcal{A}(H)$ also embeds into $\mathcal{A}(K)$. Recall that $\mathcal{A}(K)$ is a permutation automaton: up to conjugating K, we may assume that the distinguished vertex of $\mathcal{A}(K)$ is 1. Then $1 \cdot x = q \neq 1$ in $\mathcal{A}(K)$, since that is the case in $\mathcal{A}(H')$. Therefore $x \notin K$, as we wanted to show.

So we have shown that if the automaton $\mathcal{A}(H)$ does not have any vertices of degree 1, then H is closed. If $\mathcal{A}(H)$ has a vertex of degree 1, it has only one, namely the vertex 1. If we iteratively remove from $\mathcal{A}(H)$ the vertices of degree 1, we delete from $\mathcal{A}(H)$ a path starting at 1. Let x be the label of that path, and let Q^- be the resulting set of states. Let H' be the subgroup of $F(A)$ such that $\mathcal{A}(H') = (Q^-, 1 \cdot x, \cdot)$. Then as above, we can show that $H = xH'x^{-1}$. Since we are assuming that all the conjugates of H are extendible, then all the conjugates of H' are extendible, and by the first part of the proof, it follows that H' is closed. But we already observed that being closed is preserved by the automorphisms of the free group, so H too is closed. $\square$

2.4 The case where $\mathbf{V}$ is extension-closed

When the pseudovariety of groups $\mathbf{V}$ is closed under extension, the different notions we have considered so far (being closed, being extendible, and being a free factor of a clopen subgroup) all coincide. This was first proved by Ribes and Zalesskiĭ [18, Corollary 3.3 and Corollary 3.8].

Proposition 2.17 *Let H be a finitely generated subgroup of $F(A)$. If $\mathbf{V}$ is extension-closed, the following are equivalent.*

(1) *H is closed.*

(2) *H is a free factor of a clopen subgroup.*

(3) *H is extendible.*

Moreover the pro-$\mathbf{V}$ topology on H coincides with the restriction to H of the topology of $F(A)$.

Proof. We already know that (1) implies (3), and (3) implies (2). To prove the last implication, let us assume that H is a free factor of a clopen subgroup K of $F(A)$. Then H is closed in the pro-$\mathbf{V}$ topology of K by Corollary 1.8. (To apply this corollary, we need to know that $F(A)$ is residually $\mathbf{V}$, but that is always the case when $\mathbf{V}$ is extension closed.) Now the pro-$\mathbf{V}$ topology of K coincides with the restriction to K of the pro-$\mathbf{V}$ topology of $F(A)$ by Proposition 1.9. So H is closed in $F(A)$ as well.

The last statement is a consequence of Proposition 1.6 and Proposition 1.9. $\square$

Corollary 2.18 *Let H be a finitely generated subgroup of $F(A)$. If $\mathbf{V}$ is extension-closed, then $\tilde{H} = Cl(H)$.*

Proof. By Proposition 2.9, $\tilde{H}$ is the least extendible subgroup containing H. By Prop 2.17, it is also the least closed subgroup containing H, and hence it is equal to $Cl(H)$. $\square$

This implies that if $\mathbf{V}$ is extension-closed, then the closure of a finitely generated subgroup also has finite rank. In fact, the following stronger result holds. (This is [18, Proposition 3.4], with a slightly more direct proof.)

Proposition 2.19 *Let H be a finitely generated subgroup of $F(A)$. If $\mathbf{V}$ is extension-closed, then $rk(Cl(H)) \leq rk(H)$.*

Proof. By Proposition 2.17, the pro-$\mathbf{V}$ topology on $Cl(H)$ coincides with the restriction to $Cl(H)$ of the pro-$\mathbf{V}$ topology on $F(A)$. In particular, H is dense in the pro-$\mathbf{V}$ topology of $Cl(H)$.

So it suffices to show that for a subgroup H of the free group $F(A)$ to be dense, the rank of H must be at least equal to $|A|$. Let $\jmath: H \to F(A)$ be the natural morphism from H into $F(A)$, and let $\hat\jmath$ be the continuous morphism from the pro-$\mathbf{V}$ completion $\hat H$ of H into $\hat F_A(\mathbf{V})$, the completion of $F(A)$. Since the free group is residually $\mathbf{V}$ (a consequence of the fact that $\mathbf{V}$ is extension-closed), the following diagram is commutative, and the natural morphisms from H and $F(A)$ into their respective completions are one-to-one.

$$\begin{tikzcd} H \arrow[r, "\jmath"] \arrow[d, "\hat\imath"'] & F(A) \arrow[d, "\imath"] \\ \hat H \arrow[r, "\hat\jmath"] & \hat F_A(\mathbf{V}) \end{tikzcd}$$

We have $\overline{\imath\jmath(H)} = \overline{\hat\jmath\imath_H(H)}$. By compactness, it follows that $\overline{\imath\jmath(H)} = \hat\jmath(\overline{\imath_H(H)}) = \hat\jmath(\hat H)$. By assumption, the closure of $\jmath(H)$ is $F(A)$, so the closure of $\imath\jmath(H)$ is $\overline{\imath\jmath(H)} = \overline{\imath(F(A))} = \hat F_A(\mathbf{V})$. Thus $\hat\jmath$ is a continuous morphism from $\hat H$ onto $\hat F_A(\mathbf{V})$. But the notion of rank of a free pro-$\mathbf{V}$ group is well defined, and the rank of $\hat H$ is that of H [6, Lemma 15.19], so the announced inequality follows immediately. $\square$

Note. As we will see, if $\mathbf{V}$ is the pseudovariety of nilpotent groups (a pseudovariety which is not extension closed), the $\mathbf{V}$-closure of a finitely generated subgroup always has finite rank (Proposition 4.3 below). However, in view of Example 2.14 above, the inequality in Proposition 2.19 does not always hold.

Finally, Proposition 2.16 implies the following.

Corollary 2.20 *Let H be a finitely generated subgroup of $F(A)$. If $\mathbf{V}$ is extension-closed, then H is extendible if and only if any one of its conjugates is extendible.*

3 Practical computation: the pro-p topology

In this section, we give a new version of Ribes and Zalesskiĭ's algorithm to compute the closure of a finitely generated subgroup in the pro-p topology (where p is a prime number) [18]. This new version of the algorithm seems to us to be a bit clearer, and it allows us to give a polynomial upper bound to the complexity of the problem of computing such a closure.

Before we embark in the description and analysis of the algorithm, let us make some simple remarks.

Lemma 3.1 *Let* $\mathbf{V}$ *and* $\mathbf{W}$ *be pseudovarieties of groups with* $\mathbf{V} \subseteq \mathbf{W}$*. Let* H *be a finitely generated subgroup of the free group* $F(A)$*. If* H *is* $\mathbf{V}$*-clopen, then it is also* $\mathbf{W}$*-clopen.*

Proof. This is immediate by Proposition 2.5. Indeed, if H is $\mathbf{V}$-clopen, then $M(H) \in \mathbf{V}$, so $M(H) \in \mathbf{W}$ and H is $\mathbf{W}$-clopen. □

Corollary 3.2 *Let* $\mathbf{V}$ *and* $\mathbf{W}$ *be pseudovarieties of groups with* $\mathbf{V} \subseteq \mathbf{W}$*. Let* H *be a finitely generated subgroup of the free group* $F(A)$*.*

1. *If* H *is* $\mathbf{W}$*-dense, then* H *is* $\mathbf{V}$*-dense.*
2. *The* $\mathbf{W}$*-closure of* H *is contained in the* $\mathbf{V}$*-closure of* H*.*
3. *If* H *is* $\mathbf{V}$*-closed, then* H *is* $\mathbf{W}$*-closed.*

Proof. If H is not $\mathbf{V}$-dense, then H is contained in a proper $\mathbf{V}$-clopen subgroup, so by Lemma 3.1, H is contained in a proper $\mathbf{W}$-clopen subgroup, and hence H is not $\mathbf{W}$-dense. This proves the first statement.

The second statement follows directly from Lemma 3.1 since the closure of H is the intersection of the clopen subgroups containing H. The last statement is an immediate consequence of the second one. □

Proposition 3.3 *Let* $\mathbf{V}$ *be a pseudovariety of groups. If, given a finitely generated subgroup* H *of* $F(A)$*, one can decide membership in* $Cl(H)$*, then one can effectively compute* $\tilde{H}$*.*

Proof. For each state q of $\mathcal{A}(H)$, fix an element $u_q \in F(A)$ such that $1 \cdot u_q = q$ in $\mathcal{A}(H)$. We want to compute the congruence $\sim$ on $\mathcal{A}(H)$, that is, the congruence $\sim_{H,Cl(H)}$. So we must decide, for each pair of distinct states r and s whether the word $u_r u_s^{-1}$ lies in $Cl(H)$. This simple observation completes the proof. □

We now turn to the specifics of the situation when $\mathbf{V}$ is the pseudovariety of p-groups. In the rest of this section, $\mathbf{V}$ is set equal to $\mathbf{G}_p$ for some prime number p. We will talk of p-closure, p-denseness, etc. instead of $\mathbf{G}_p$-closure, $\mathbf{G}_p$-denseness, etc.

3.1 Deciding p-denseness

We use the following key property: in a finite p-group, every maximal proper subgroup is normal of index p [19, Thm 4.6].

Lemma 3.4 *If H is a proper p-clopen subgroup of $F(A)$, then there exists an onto morphism $\psi\colon F(A) \to \mathbb{Z}/p\mathbb{Z}$ such that $H \subseteq \ker\psi$.*

Proof. Recall that if $\mu_H\colon F(A) \to M(H)$ is the transition morphism of $\mathcal{A}(H)$, then $H = \mu_H^{-1}\mu_H(H)$. In particular, $\mu_H(H) \neq M(H)$ since H is proper. Let N be a maximal proper subgroup of $M(H)$ containing $\mu_H(H)$. Then N is normal and has index p. Let π be the projection from $M(H)$ onto $M(H)/N = \mathbb{Z}/p\mathbb{Z}$, and let $\psi = \pi \circ \varphi\colon F(A) \to \mathbb{Z}/p\mathbb{Z}$. Then $\ker\psi = \varphi^{-1}(N)$ contains H. $\square$

The pseudovariety $(\mathbb{Z}/p\mathbb{Z})$ generated by $\mathbb{Z}/p\mathbb{Z}$ admits a finite free object over the alphabet A, namely $(\mathbb{Z}/p\mathbb{Z})^A$. Let $\sigma\colon F(A) \to (\mathbb{Z}/p\mathbb{Z})^A$ be the natural onto morphism.

Lemma 3.4, together with Proposition 2.2, immediately implies the following.

Corollary 3.5 *Let H be a finitely generated subgroup of $F(A)$. The following are equivalent.*

1. *H is p-dense.*
2. *H is $(\mathbb{Z}/p\mathbb{Z})$-dense.*
3. *$\sigma^{-1}\sigma(H) = F(A)$.*
4. *$\sigma(H) = (\mathbb{Z}/p\mathbb{Z})^A$.*

The problem of determining whether the subgroup H is p-dense can now be turned into a question of linear algebra over the p-element field: we need to determine the dimension of the subspace $\sigma(H)$ in $(\mathbb{Z}/p\mathbb{Z})^A$.

Let $h_1, \ldots, h_r$ be a given set of generators for H. Let $\mathfrak{M}_p(H)$ be the $r \times |A|$ matrix consisting of the row vectors $\sigma(h_1), \ldots, \sigma(h_r)$, a generating set for $\sigma(H)$.

Corollary 3.6 *H is p-dense if and only if $\mathfrak{M}_p(H)$ has rank $|A|$.*

Note. It is clear that if $r < |A|$, then H is not p-dense. Naturally, this is exactly the proof that the rank of free pro-$\mathbf{V}$ groups ($\mathbf{V}$ non-trivial) is uniquely defined.

We can compute the rank of the matrix $\mathfrak{M}_p(H)$ by Gaussian elimination: by taking linear combinations of the rows, we obtain an upper triangular matrix T such that the number of non-zero rows of T is the rank of $\mathfrak{M}_p(H)$, and the non-zero rows of T form a basis of $\sigma(H)$. This computation can be done in time $O(r^2|A|)$.

Corollary 3.7 *It is decidable whether H is p-dense. In addition, we can compute a basis of $\sigma(H)$ if H is not p-dense.*

3.2 Computing the p-closure

Let H be a finitely generated subgroup of $F(A)$. We assume that we are given a basis (of reduced words) for H, the automaton $\mathcal{A}(H)$, and for each state q of $\mathcal{A}(H)$, a reduced word u_q which labels a path from 1 to q in $\mathcal{A}(H)$. We can assume that the length of each u_q is less than or equal to the number of states of $\mathcal{A}(H)$. See Section 2.1 on how to compute $\mathcal{A}(H)$ from a basis of H and vice versa.

To compute the closure of H, we compute a finite sequence of quotients of $\mathcal{A}(H)$, $\mathcal{A}(H_0) = \mathcal{A}(H)/\sim_0, \ldots, \mathcal{A}(H_n) = \mathcal{A}(H)/\sim_n$, such that each H_i is p-closed, the automaton congruence $\sim_{i+1}$ is contained in $\sim_i$ (that is $H_{i+1} \subseteq H_i$), and H_n is the p-closure of H. Since each H_i is p-closed, each contains the p-closure of H. Moreover, it follows from Proposition 2.17 that the p-closure of H in $F(A)$ is equal to the p-closure of H in the pro-p topology of H_i.

To begin with, we let $\sim_0$ be the universal, one-class congruence, so that $H_0 = F(A)$. Let $i \geq 0$. After i iterations of the algorithm, we have computed the quotient $\mathcal{A}(H_i) = \mathcal{A}(H)/\sim_i$. Roughly speaking, for the $(i+1)$-st iteration of the algorithm, we translate H into a basis of H_i and we ask whether H is p-dense in H_i. If it is, H_i is the closure of H; if not, we compute the $(\mathbb{Z}/p\mathbb{Z})$-closure of H in H_i, or rather a free factor H_{i+1} of that closure which contains H. Formally, we proceed as follows.

[1] *Computing a basis of H_i.* First we compute a basis for H_i. Let A_i be a set in bijection with that basis. We let $\kappa_i \colon F(A_i) \to H_i \subseteq F(A)$ be the natural one-to-one morphism onto H_i. More precisely, we choose a spanning tree of $\mathcal{A}(H_i)$: then A_i is (in bijection with) the set of edges of $\mathcal{A}(H_i)$ not in that spanning tree (see Section 2.1). We do not, in fact, need to write down explicitly the elements of the basis of H_i, that is, the $\kappa_i(x)$ $(x \in A_i)$. Note that κ_i is a homeomorphism between $F(A_i)$ and H_i. We denote by σ_i the natural morphism $\sigma_i\colon F(A_i) \to (\mathbb{Z}/p\mathbb{Z})^{A_i}$.

$$\begin{tikzcd} F(A_i) \arrow[r, "\sigma_i"] \arrow[d, "\kappa_i"'] & (\mathbb{Z}/p\mathbb{Z})^{A_i} \\ H_i & \end{tikzcd}$$

[2] *Translating H into the basis of H_i.* Now we compute a basis of the subgroup $\kappa_i^{-1}(H)$ of $F(A_i)$ (this is equivalent to rewriting the basis of H in the chosen basis of H_i). This is done by running the elements of the basis of H in $\mathcal{A}(H_i)$ and noting down the edges traversed that are not in the chosen spanning tree. This set of words over the alphabet $A_i \cup A_i^{-1}$ is indeed a basis since κ_i is a homeomorphism onto its image.

[3] *Deciding the p-denseness of H in H_i.* Then we use the algorithm in Section 3.1 to decide whether $\kappa_i^{-1}(H)$ is p-dense in $F(A_i)$, and to compute a basis of $\sigma_i(\kappa_i^{-1}(H))$ if it is not p-dense. Observe that H is p-dense in H_i if and only if $\kappa_i^{-1}(H)$ is p-dense in $F(A_i)$ since κ_i is a homeomorphism.

[4] *Stop if H is p-dense in H_i.* If H is p-dense in H_i, the algorithm stops: we now know that the p-closure of H (in H_i, and hence in $F(A)$) is H_i.

[5] *Otherwise compute H_{i+1}.* We now assume that $\kappa_i^{-1}(H)$ is not p-dense in $F(A_i)$. By Proposition 2.2, $\sigma_i^{-1}\sigma_i\kappa_i^{-1}(H)$ is the $(\mathbb{Z}/p\mathbb{Z})$-closure of $\kappa_i^{-1}(H)$ in $F(A_i)$, and by Corollary 3.5 it is properly contained in $F(A_i)$. Since κ_i is a homeomorphism from $F(A_i)$ onto H_i, the subgroup $K = \kappa_i\sigma_i^{-1}\sigma_i\kappa_i^{-1}(H)$ is the $(\mathbb{Z}/p\mathbb{Z})$-closure of H in H_i and $K \neq H_i$. Finally, by Corollary 3.2, K is p-closed in the pro-p topology of H_i, and hence also in $F(A)$ (Proposition 2.17). We define the automaton congruence $\sim_{i+1}$ on $\mathcal{A}(H)$ to be $\sim_{H,K}$, the congruence induced by the containment of H into K. In particular, the subgroup H_{i+1} such that $\mathcal{A}(H_{i+1}) = \mathcal{A}(H)/\sim_{i+1}$ is a free factor of K, and hence H_{i+1} is p-closed. Moreover, we have $H \subseteq H_{i+1} \subseteq K \subset H_i$, and hence H_{i+1} is properly contained in H_i and $\sim_{i+1}$ is properly contained in $\sim_i$.

The automaton congruence $\sim_{i+1}$ is computed as follows. If r and s are states of $\mathcal{A}(H)$, we have $r \sim_{i+1} s$ if and only if $u_r u_s^{-1} \in K = \kappa_i\sigma_i^{-1}\sigma_i\kappa_i^{-1}(H)$, that is, if and only if $u_r u_s^{-1} \in H_i$ and $\sigma_i \kappa_i^{-1}(u_r u_s^{-1}) \in \sigma_i \kappa_i^{-1}(H)$. To verify whether $u_r u_s^{-1} \in H_i$, and to compute in that case $\kappa_i^{-1}(u_r u_s^{-1})$, we run the reduced word obtained from $u_r u_s^{-1}$ in the automaton $\mathcal{A}(H_i)$ starting at 1, we note down the edges traversed that are not in the chosen spanning tree of that automaton (as in step 2), and we require that this path ends in 1. Then $\sigma_i \kappa_i^{-1}(u_r u_s^{-1})$ is the image of that word in $(\mathbb{Z}/p\mathbb{Z})^A$. Now it suffices to verify whether the vector $\sigma_i \kappa_i^{-1}(u_r u_s^{-1})$ lies in the vector subspace $\sigma_i \kappa_i^{-1}(H)$. This can be done effectively, using the basis of $\sigma_i \kappa_i^{-1}(H)$ computed in step 3.

Finally we observe that the algorithm stops after $i+1$ iteration exactly if H_i is the p-closure of H (see step 2).

Example 3.8 In Example 1.10, we considered the subgroups

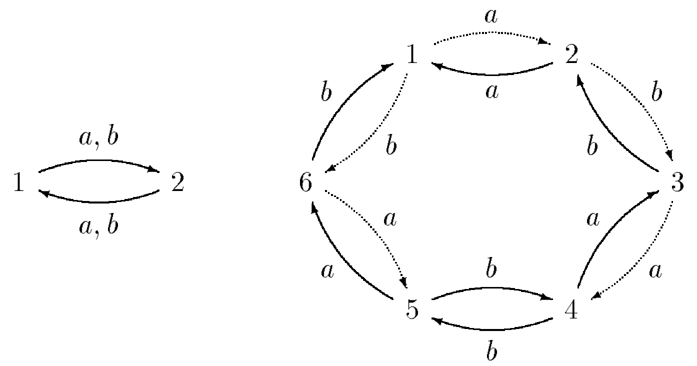

$H = \langle a^2, ab, ab^{-1} \rangle$ $\qquad$ $K = \langle a^2, ab^2a^{-1}, aba^2b^{-1}a^{-1}, ababa^{-1}b^{-1}, baba^{-1}b^{-1}a^{-1}, ba^2b^{-1}, b^2 \rangle$

A spanning tree of $\mathcal{A}(K)$ is indicated by the dotted arrows. In order to compute the p-closure of K, we first consider the image of $\sigma_0(K)$ of K in $(\mathbb{Z}/p\mathbb{Z})^{\{a,b\}}$: it is generated by $2a$, $2b$, $a+b$ and $b-a$. We need to compute the rank of the matrix

$$\begin{pmatrix} 2 & 0 \\ 0 & 2 \\ 1 & 1 \\ -1 & 1 \end{pmatrix}$$

For $p \neq 2$, this rank is easily seen to be 2, so K is p-dense in $F(A)$. For $p = 2$, the rank of the matrix is 1, and $\sigma_0(K)$ is the subspace generated by $a+b$. For the given choice of a spanning tree, we have the following values of the words u_q:

q	1	2	3	4	5	6
u_q	1	a	ab	aba	ba	b

Now $\sigma_0(u_1u_3^{-1}) = \sigma_0(u_1u_5^{-1}) = a + b$, so $1 \sim_1 3 \sim_1 5$, and similarly, $2 \sim_1 4 \sim_1 6$. It follows that $\mathcal{A}(H_1) = \mathcal{A}(K)/\sim_1 = \mathcal{A}(H)$, so $H_1 = H$.

Let us denote by x, y and z respectively the elements a^2, ab and ba^{-1} of the basis of H. Then

$$\kappa_1^{-1}(K) = \langle x, yz, yxy^{-1}, y^2x^{-1}z^{-1}, zxzy^{-1}, zxz^{-1}, zy \rangle.$$

The resulting matrix (over $\mathbb{Z}/2\mathbb{Z}$) is

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 1 & 0 & 0 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 1 \end{pmatrix}$$

which has rank 3. Thus $\kappa_1^{-1}(K)$ is dense in $F(\{x, y, z\})$, that is, K is 2-dense in H, and hence H is the 2-closure of K.

Example 3.9 Let H be the group of Example 2.14:

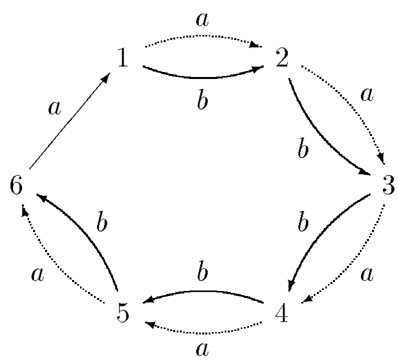

$$H = \langle ab^{-1}, a^2b^{-1}a^{-1}, a^3b^{-1}a^{-2}, \\ a^4b^{-1}a^{-3}, a^5b^{-1}a^{-4}, a^6 \rangle$$

A spanning tree is indicated by dotted arrows. The matrix whose rank we must compute is

$$\begin{pmatrix} 1 & -1 \\ 6 & 0 \end{pmatrix}$$

For $p \geq 5$, this matrix has rank 2, so H is p-dense in $F(A)$.

For $p = 2$, $\sigma_0(H)$ is generated by $a + b$. We easily find that $1 \sim_1 3 \sim_1 5$ and $2 \sim_1 4 \sim_1 6$. So the automaton of H_1 is

$$\begin{tikzcd} \{1,3,5\} \arrow[r, bend left, "{a,b}"] & \{2,4,6\} \arrow[l, bend left, "{a,b}"] \end{tikzcd} \qquad H_1 = \langle ab^{-1}, a^2, ab \rangle$$

Then one verifies that H is 2-dense in H_1, that is, H_1 is the 2-closure of H.

For $p = 3$, $\sigma_0(H)$ is generated by $a - b$. Then one can verify that $1 \sim_1 4$, $2 \sim_1 5$ and $3 \sim_1 6$, so the first iteration of the algorithm of computation of the 3-closure of H yields the following automaton and subgroup:

$$\begin{tikzcd} \{1,4\} \arrow[rr, "{a,b}"] & & \{2,5\} \arrow[dl, "{a,b}"] \\ & \{3,6\} \arrow[ul, "{a,b}"] & \end{tikzcd} \qquad H_1' = \langle ba^{-1}, aba^{-2}, a^2b, a^3 \rangle$$

One then verifies that H is 3-dense in H_1', so H_1' is the 3-closure of H.

The last example shows a less simple case, where the algorithm needs to undergo several iterations before it stops.

Example 3.10 Let us compute the 2-closure of the following group:

$$\begin{aligned} H \;=\; \langle\, & a^4, a^2b^{-2}, ab^2a, ba^4b^{-1}, ba^2b^{-3}, bab^{-2}ab^{-1} \\ & aba^4b^{-1}a^{-1}, aba^2b^{-3}a^{-1}, abab^2ab^{-1}a^{-1}, \\ & bababa^{-1}ba, baba^{-1}baba, bab^{-1}ababa, bab^{-1}a^{-1}ba^{-1}ba \,\rangle. \end{aligned}$$

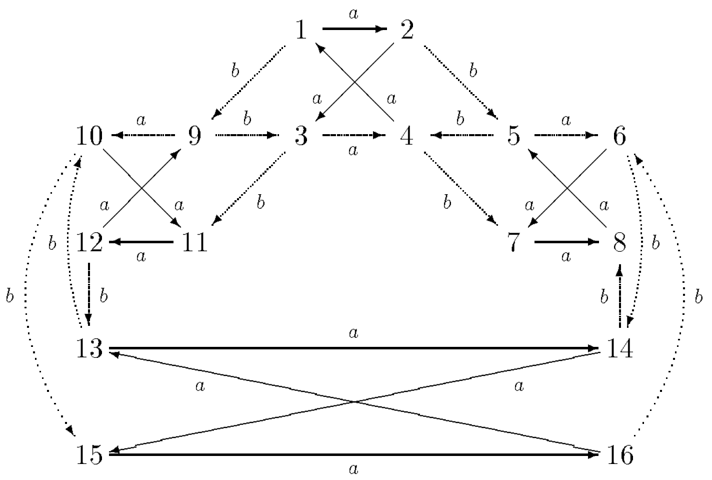

The dotted arrows indicate a spanning tree. For each state q, we compute the label u_q of a geodesic from 1 to q in the spanning tree:

q	1	2	3	4	5	6
u_q	1	b^2ab^{-2}	b^2	b^2a	b^2ab^{-1}	$b^2ab^{-1}a$
q	7	8	9	10	11	12
u_q	b^2ab	$b^2ab^{-1}ab^2$	b	ba	b^3	bab^{-2}
q	13	14	15	16		
u_q	bab^{-1}	$b^2ab^{-1}ab$	bab	$b^2ab^{-1}ab^{-1}$		

We have $A = A_0 = \{a, b\}$, H_0 is the free group $F(A)$, κ_0 is the identity of $F(A)$, and σ_0 is the natural morphism from $F(A)$ onto $(\mathbb{Z}/2\mathbb{Z})^A$. Then $\sigma_0(H) = 0$.

We compute for which pairs of states (q, r) we have $\sigma_0(u_q u_r^{-1}) = 0$, and we find that $\sim_1$ has 4 classes:

$$1 \sim_1 3 \sim_1 14 \sim_1 16 \qquad 2 \sim_1 4 \sim_1 13 \sim_1 15$$
$$5 \sim_1 7 \sim_1 10 \sim_1 12 \qquad 6 \sim_1 8 \sim_1 9 \sim_1 11.$$

So $\mathcal{A}(H_1)$ is the following automaton:

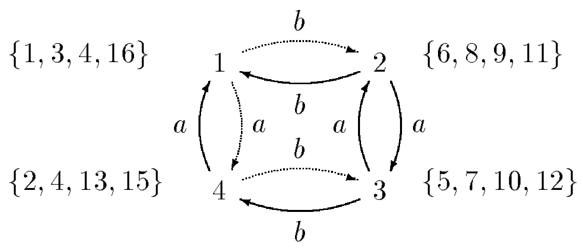

where state 1 corresponds to the class $\{1, 3, 14, 16\}$, state 2 corresponds to $\{6, 8, 9, 11\}$, etc. A spanning tree is indicated with dotted lines. The element of the induced basis of H_1 corresponding to the edge labeled a (resp. b) out of vertex i is denoted a_i (resp. b_i). Thus, H_1 has basis a_2, a_3, a_4, b_2, b_3 (e.g. $a_2 = bab^{-1}a^{-1}$), and

$$\begin{aligned}\kappa_1^{-1}(H) \quad = \quad & \langle a_4^2, a_4b_2^{-1}, b_3a_4, (a_2a_3)^2, a_2a_3b_2^{-1}, a_2b_3^{-1}a_3 \\ & (a_3a_2)^2, a_3a_2b_3^{-1}, a_3b_2a_2 \\ & a_2b_3a_4a_3^{-1}b_3a_4, a_2b_3a_2b_3a_4, a_2a_4a_2b_3a_4, a_2a_3^{-1}b_3a_4\rangle.\end{aligned}$$

Then one verifies that $\sigma_1\kappa_1^{-1}(H)$ is the vector subspace generated by $\{a_2 + a_3, a_4, b_2, b_3\}$, that is, the set of vectors in which the a_2- and a_3-components are equal. Now we compute $\sigma_1\kappa_1^{-1}(u_q u_r^{-1})$ for each pair of states (q, r) such that $q \sim_1 r$.

For instance, we find that $u_1u_3^{-1} = b^{-2}$, so $\kappa_1^{-1}(u_1u_3^{-1}) = b_2^{-1}$, and $\sigma_1\kappa_1^{-1}(u_1u_3^{-1}) = b_2 \in \sigma_1\kappa_1^{-1}(H)$. Thus $1 \sim_2 3$. In contrast, $u_1u_{14}^{-1} = b^{-1}a^{-1}ba^{-1}b^{-2}$, $\kappa_1^{-1}(u_1u_{14}^{-1}) = b_2^{-1}a_3^{-1}b_3b_2^{-1}$, and $\sigma_1\kappa_1^{-1}(u_1u_{14}^{-1}) = a_3 + b_3 \notin \sigma_1\kappa_1^{-1}(H)$. Thus $1 \not\sim_2 14$.

The complete computation yields 8 classes:

$$\begin{array}{llll} 1 \sim_2 3 & 14 \sim_2 16 & 2 \sim_2 4 & 13 \sim_2 15 \\ 5 \sim_2 7 & 10 \sim_2 12 & 6 \sim_2 8 & 9 \sim_2 11. \end{array}$$

So $\mathcal{A}(H_2)$ is the following automaton:

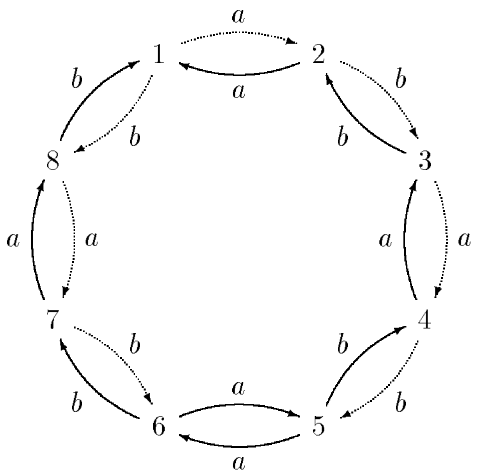

where states 1 through 8 correspond respectively to the classes $\{1, 3\}$, $\{2, 4\}$, $\{5, 7\}$, $\{6, 8\}$, $\{14, 16\}$, $\{13, 15\}$, $\{10, 12\}$ and $\{9, 11\}$. A spanning tree is indicated with dotted lines. With the same convention as before, the elements of the induced basis of H_2 are $a_2, a_4, a_5, a_6, a_7, b_3, b_5, b_6, b_8$ (e.g. $a_2 = a^2$ and $a_4 = aba^2b^{-1}a^{-1}$), and

$$\begin{array}{rcl} \kappa_2^{-1}(H) & = & \langle a_2^2, a_2b_8^{-1}, b_3a_2, a_7^2, a_7b_8^{-1}, b_6^{-1}a_7 \\ & & a_4^2, a_4b_3^{-1}, b_5a_4 \\ & & a_6b_5b_3a_2, a_5^{-1}b_5a_4b_3a_2, b_6^{-1}a_6b_5a_4b_3a_2, b_6^{-1}a_5^{-1}b_5b_3a_2\rangle. \end{array}$$

Then one verifies that no pair of distinct $\sim_2$-equivalent states are $\sim_3$-equivalent. That is, $\sim_3$ is the identity relation, and hence $H_3 = H$. H is trivially dense in H_3, so the algorithm stops, and we have shown that H is 2-closed.

3.3 Complexity issues

We consider the time complexity of each iteration of the algorithm of the previous section. We let n be the total length of the initial input, namely the sum of the lengths of the elements of the basis of H. We let also q be the number of states of $\mathcal{A}(H)$ and $a = |A|$. Then $q \leq n$.

After i iterations of the automaton, we have computed $\mathcal{A}(H_i)$. Let q_i (resp. e_i) be the number of states (resp. edges) of $\mathcal{A}(H_i)$. We let $a_0 = e_0 = a$.

In step 1 of the $(i+1)$-st iteration of the algorithm, we compute a spanning tree of $\mathcal{A}(H_i)$. This is done in time $O(e_i)$. Let a_i be the cardinality of a basis of H_i, that is, the number of edges of $\mathcal{A}(H_i)$ not in the spanning tree. Then $a_i = e_i - q_i + 1$.

Step 2, where we compute a basis of $\kappa_i^{-1}(H)$, is performed by running the elements of the basis of H in $\mathcal{A}(H_i)$. When each word is run (from 1 to 1), we store the list of edges of $\mathcal{A}(H_i)$ not in T traversed. The tree T contains $q_i - 1$ edges. Thus step 2 can be performed in time $O(nq_i)$. The sum of the lengths of the words thus obtained, say n_i, is therefore less than or equal to n.

Step 3 is performed in time $O(r^2 a_i)$, where r is the rank of H, by the results of Section 3.1. Observe that $r \leq n$.

Step 4 takes constant time.

Step 5. To perform step 5, it suffices to compute, for each pair of distinct states (r, s) of $\mathcal{A}(H)$, whether $u_r u_s^{-1}$ lies in H_i, and if so, whether $\sigma_i \kappa_i^{-1}(u_r u_s^{-1})$ lies in $\sigma_i \kappa_i^{-1}(H)$. The word $u_r u_s^{-1}$ is of length at most $2q$, reducing it takes time $O(q)$, then verifying whether it lies in H_i and if so, computing $\sigma_i \kappa_i^{-1}(u_r u_s^{-1})$ is done in time $O(qq_i)$ as in step 2. Next, computing whether $\sigma_i \kappa_i^{-1}(u_r u_s^{-1})$ lies in $\sigma_i \kappa_i^{-1}(H)$ can be done in time $O(na_i)$ (we have already computed a basis of $\sigma_i \kappa_i^{-1}(H)$ in step 3). Thus step 5 can be performed in time $O(q^2(qq_i + na_i))$.

Thus the $(i+1)$st iteration of the algorithm takes time $O(q^3 q_i + q^2 n a_i)$.

To evaluate the time complexity of the complete algorithm, we remark that $a_i \leq e_i \leq n$ and $q_i \leq e_i$ for each i. So each iteration of the algorithm takes time $O(n^4)$. Finally, for each i, $\sim_{i+1}$ is strictly contained in $\sim_i$, so $q_i < q_{i+1} \leq q$. So the algorithm will stop after at most q iterations. But we remarked that $q \leq n$, so we have the following statement.

Theorem 3.11 *Let p be a prime number. If H is a finitely generated subgroup of the free group $F(A)$, given by a finite set of generators, and if n*

is the sum of the lengths of these given generators, then we can construct the automata $\mathcal{A}(Cl(H))$ and $\mathcal{A}(\tilde{H})$ in time $O(n^5)$ (with respect to the pro-p topology).

Note. The algorithm discussed in this and the previous section can be used to solve the following problem: given m permutations on an n-element set, decide whether the group generated by these permutations is a p-group. It suffices to consider the inverse automaton representing the action of the m permutations (now seen as m letters) on the n-element set, and to decide whether the corresponding subgroup of the rank m free group (a finite index subgroup) is p-closed. However this decision procedure has a time complexity bounded above by $O(m^5n^5)$, whereas the problem is known to be solvable in time $O(mn\log n\log^* n)$ [15].

Let us conclude this section with a simple remark. If an overgroup K of H (i.e. $\mathcal{A}(K)$ is a quotient of $\mathcal{A}(H)$) properly contained in $F(A)$ has rank greater than $rk(H)$, then it cannot be the p-closure of H for any prime p (by Proposition 2.19). In view of Corollary 3.6 this yields the following proposition.

Proposition 3.12 *Let H be a finitely generated subgroup of the free group $F(A)$ such that every overgroup of H properly contained in $F(A)$ has rank greater than $rk(H)$. Let p be a prime number. Then H is p-dense if and only if the matrix $\mathfrak{M}_p(H)$ has rank $|A|$, and H is p-closed otherwise.*

This result applies to what we could call *primitive* subgroups of the free group, that is, subgroups H such that $\mathcal{A}(H)$ admits no non-trivial congruence. It would be interesting to have more information on the structure of such subgroups.

3.4 Computation of extensions

We have seen that a finitely generated subgroup H of the free group $F(A)$ is p-closed if and only if the automaton $\mathcal{A}(H)$ can be embedded into the automaton $\mathcal{A}(K)$ of a p-clopen subgroup K, that is, into a permutation automaton whose transition group is a p-group. We now explain how to effectively construct $\mathcal{A}(K)$ when H is given. The procedure is an extension of that of Section 3.2.

Given a subgroup H of $F(A)$, we have constructed in Section 3.2 a finite sequence of subgroups $H_0 = F(A)$, H_1, ..., H_r such that $H_r = Cl_p(H)$.

We now construct in addition p-clopen subgroups $K_0, K_1, \ldots, K_r$ such that $\mathcal{A}(H_i)$ is embedded in $\mathcal{A}(K_i)$ for each $0 \leq i \leq r$. Since $H_0 = F(A)$, we let $K_0 = F(A)$ as well.

Now suppose that we have computed H_i (as in Section 3.2) and K_i ($i \geq 0$) such that K_i is p-clopen and $\mathcal{A}(H_i)$ embeds into $\mathcal{A}(K_i)$. We use the notation of Section 3.2.

In step 1 of the $(i+1)$-st iteration we compute not only a basis of H_i, but also a basis of K_i which contains a basis of H_i. We let B_i be a set containing A_i, equipped with a bijective correspondence with the basis of K_i which extends the bijection between A_i and the basis of H_i. More precisely, we choose a spanning tree of $\mathcal{A}(H_i)$, and then extend it to a spanning tree of $\mathcal{A}(K_i)$: then B_i is the set of edges of $\mathcal{A}(K_i)$ not in the tree, and A_i is the set of edges of $\mathcal{A}(H_i)$ not in the tree.

Let $\lambda_i\colon F(B_i) \to K_i$ be the natural isomorphism. The restriction of λ_i to $F(A_i)$ is κ_i. Let also $\tau_i\colon F(B_i) \to (\mathbb{Z}/p\mathbb{Z})^{B_i}$ be the natural morphism, and let $\pi_i\colon (\mathbb{Z}/p\mathbb{Z})^{B_i} \to (\mathbb{Z}/p\mathbb{Z})^{A_i}$ be the projection onto the A_i-coordinates. The following diagram is commutative (where μ_i and ν_i are the natural injections):

$$\begin{tikzcd}
K_i & F(B_i) \arrow[l, "\lambda_i"'] \arrow[r, "\tau_i"] & (\mathbb{Z}/p\mathbb{Z})^{B_i} \arrow[d, "\pi_i"] \\
H_i \arrow[u, "\mu_i"] & F(A_i) \arrow[l, "\kappa_i"'] \arrow[u, "\nu_i"] \arrow[r, "\sigma_i"] & (\mathbb{Z}/p\mathbb{Z})^{A_i}
\end{tikzcd}$$

Step 2 of the $(i+1)$-st iteration consists in computing (a basis of) $\lambda_i^{-1}(H)$. But $H \subseteq H_i \subseteq K_i$, so $\lambda_i^{-1}(H) = \kappa_i^{-1}(H)$, and this step is identical to the corresponding step in Section 3.2. Steps 3 and 4 also are unchanged from Section 3.2.

In step 5 we compute H_{i+1} as in Section 3.2, and we also compute K_{i+1}: we let $K_{i+1} = \lambda_i\tau_i^{-1}\tau_i\lambda_i^{-1}(H)$. Then K_{i+1} is the $(\mathbb{Z}/p\mathbb{Z})$-closure of H in K_i, whereas H_{i+1} is a free factor of $\kappa_i\sigma_i^{-1}\sigma_i\kappa_i^{-1}(H)$, the $(\mathbb{Z}/p\mathbb{Z})$-closure of H in H_i. In particular, $H_{i+1} \subseteq K_{i+1}$ and K_{i+1} is p-clopen.

We now verify that $\kappa_i\sigma_i^{-1}\sigma_i\kappa_i^{-1}(H)$ is a free factor of $\lambda_i\tau_i^{-1}\tau_i\lambda_i^{-1}(H)$. First we observe that

$$\begin{aligned}
\kappa_i\sigma_i^{-1}\sigma_i\kappa_i^{-1}(H) &= \kappa_i(\tau_i^{-1}\pi_i^{-1}\pi_i\tau_i\kappa_i^{-1}(H)) \\
&= \kappa_i\tau_i^{-1}\tau_i\kappa_i^{-1}(H).
\end{aligned}$$

Next we note that $\tau_i^{-1}\tau_i(F(A_i))$ is the free product

$$\tau_i^{-1}\tau_i(F(A_i)) = F(A_i) * \tau_i^{-1}(\ker \pi_i).$$

It follows that

$$\tau_i^{-1}\tau_i\kappa_i^{-1}(H) \;=\; \sigma_i^{-1}\sigma_i\kappa_i^{-1}(H) * \tau_i^{-1}(\ker\pi_i).$$

Since λ_i is an isomorphism, we have

$$\lambda_i\tau_i^{-1}\tau_i\kappa_i^{-1}(H) \;=\; \lambda_i\sigma_i^{-1}\sigma_i\kappa_i^{-1}(H) * \lambda_i\tau_i^{-1}(\ker\pi_i).$$

But $\sigma_i^{-1}\sigma_i\kappa_i^{-1}(H) \subseteq F(A_i)$, so $\lambda_i\sigma_i^{-1}\sigma_i\kappa_i^{-1}(H) = \kappa_i\sigma_i^{-1}\sigma_i\kappa_i^{-1}(H)$. Moreover $\lambda_i^{-1}(H) = \kappa_i^{-1}(H)$, so H_{i+1} is a free factor of K_{i+1}.

So we have proved the following theorem.

Theorem 3.13 *Let H be a finitely generated closed subgroup of the free group and let p be a prime number. Then we can effectively construct a p-clopen subgroup K such that the automaton $\mathcal{A}(H)$ embeds in $\mathcal{A}(K)$.*

Note that $\tau_i\kappa_i^{-1}(H)$ is normal in $(\mathbb{Z}/p\mathbb{Z})^{B_i}$ (since the latter group is abelian), so $\tau_i^{-1}\tau_i\kappa_i^{-1}(H)$ is normal in $F(B_i)$ and K_{i+1} is normal in K_i. Moreover the quotient group K_i/K_{i+1} is isomorphic to $F(B_i)/\tau_i^{-1}\tau_i\kappa_i^{-1}(H)$, so it is a quotient of $(\mathbb{Z}/p\mathbb{Z})^{B_i}$ and hence an abelian p-group. It follows that if the algorithm to compute the p-closure of H terminates after i iterations, then the length of a decomposition tower involving only cyclic quotients, of the p-group $M(K)$ constructed as above is at most i.

4 Practical computation: the pro-nil topology

If $\mathbf{V}$ is the pseudovariety $\mathbf{G}_{nil}$ of nilpotent groups, we talk of nil-closure, nil-denseness, etc. instead of $\mathbf{G}_{nil}$-closure, $\mathbf{G}_{nil}$-denseness, etc.

4.1 General statements

Recall that $\mathbf{G}_{nil}$ is the join of the $\mathbf{G}_p$, and that, in addition, every finite nilpotent group is (isomorphic to) the direct product of its Sylow subgroups (each of which is a p-group for some prime p) [19, Thm 5.39].

Lemma 4.1 *Let H be a nil-clopen subgroup. There exists a finite collection of prime numbers $p_1, \ldots, p_n$ and subgroups $H_1, \ldots, H_n$ such that for each i, H_i is p_i-clopen, and $H = H_1 \cap \cdots \cap H_n$.*

Proof. Let $\varphi: F(A) \to M(H)$ be the transition morphism of $\mathcal{A}(H)$. Then we know that $H = \varphi^{-1}\varphi(H)$. Moreover, $M(H)$ is nilpotent, so $M(H)$ is the direct product of its Sylow subgroups, $M(H) = M_1 \times \cdots \times M_r$. For each $1 \leq i \leq r$, let p_i be the prime number such that M_i is the p_i-Sylow subgroup of $M(H)$: then M_i is the set of all elements of $M(H)$ whose order is a power of p_i. Let also $\pi_i: M(H) \to M_i$ be the i-th coordinate projection.

For each $1 \leq i \leq r$, let $N_i = \varphi(H) \cap M_i$. Then N_i is the set of elements of $\varphi(H)$ whose order is a power of p_i, that is, N_i is the p_i-Sylow subgroup of $\varphi(H)$ if p_i divides the order of $\varphi(H)$, and N_i is the trivial subgroup otherwise. Since $\varphi(H)$ is nilpotent, we have $\varphi(H) = N_1 \times \cdots N_r$. Therefore $\bigcap_{i=1}^r \pi_i^{-1}(N_i) = \varphi(H)$, and hence

$$H = \bigcap_{i=1}^{r} \varphi^{-1}\pi_i^{-1}(N_i).$$

But each $\varphi^{-1}\pi_i^{-1}(N_i)$ is a p_i-clopen subgroup, so the lemma is proved. $\square$

Corollary 4.2 *Let H be a finitely generated subgroup of the free group $F(A)$.*

1. *H is nil-dense if and only if H is p-dense for each prime p.*
2. *The nil-closure of H is the intersection over all primes p of the p-closures of H.*
3. *If H is p-closed for some p, then H is nil-closed.*

Proof. By Corollary 3.2, we already know that every nil-dense subgroup is p-dense, and that the nil-closure of H is contained in the intersection of the p-closures of H. We now show the converse statements.

Suppose that H is not nil-dense. Then H is contained in some nil-clopen proper subgroup K. By Lemma 4.1, there exist primes $p_1, \ldots, p_r$ and subgroups $K_1, \ldots, K_r$ such that K_i is p_i-clopen and $K = K_1 \cap \ldots \cap K_r$. Since K is proper, there exists $1 \leq i \leq r$ such that K_i is a proper subgroup. But H is contained in K_i, so H is not p_i-dense. This concludes the proof of the first statement.

We know that the nil-closure of H is the intersection of all the nil-clopen subgroups containing H. Let K be such a subgroup. Then as above, $K = K_1 \cap \ldots \cap K_r$, where each K_i is p_i-clopen for some prime p_i. It follows that the p_i-closure of H is contained in K_i for each $1 \leq i \leq r$. Therefore the intersection of all the p-closures of H (for all primes p) is contained in K, and hence in the nil-closure of H, thus proving the second statement. The last statement follows immediately. $\square$

Of course, there exist nil-closed subgroups which are not p-closed for any p. It suffices to consider, on one generator a, the subgroup generated by a^6. It is p-dense for $p \neq 2, 3$, its 2-closure is generated by a^2, its 3-closure is generated by a^3, and it is easily seen to be nil-clopen.

Corollary 4.2 yields the following result (see Proposition 2.19).

Proposition 4.3 *If H is a finitely generated subgroup of the free group $F(A)$, then the nil-closure of H is finitely generated. In particular, the nil-closure of H is extendible.*

Proof. For each prime p, $\mathcal{A}(Cl_p(H))$ is a quotient of $\mathcal{A}$, so there are finitely many values of $Cl_p(H)$, each of which has finite rank. So the nil-closure of H is the intersection of a finite family of finitely generated subgroups, and hence it is also finitely generated. It follows from Corollary 2.13 that the nil-closure of H is $\mathbf{G}_{nil}$-extendible. $\square$

Example 4.4 Let H and K be the groups of Examples 1.10 and 3.8:

$$\begin{aligned} H &= \langle a^2, ab, ab^{-1} \rangle \\ K &= \langle a^2, ab^2a^{-1}, aba^2b^{-1}a^{-1}, ababa^{-1}b^{-1}, \\ &\qquad baba^{-1}b^{-1}a^{-1}, ba^2b^{-1}, b^2 \rangle. \end{aligned}$$

We have seen in Example 3.8 that the p-closure of K is $F(A)$ for all values of p except for 2, and that the 2-closure of K is H. Therefore the nil-closure of K is H.

Example 4.5 Let us compute the nil-closure of the subgroup H of Examples 2.14 and 3.9:

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

In Example 3.9, we showed that the p-closure of H is $F(A)$ itself if $p \neq 2, 3$, $\langle a^2, ab, ab^{-1} \rangle$ if $p = 2$ and $\langle ba^{-1}, aba^{-2}, a^2b, a^3 \rangle$ if $p = 3$. Thus the nil-closure of H is the intersection of the above two subgroups. A simple computation shows that it is the subgroup K:

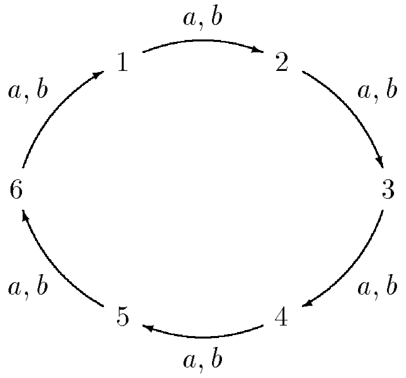

$$K = \langle ab^{-1}, a^2b^{-1}a^{-1}, a^3b^{-1}a^{-2},\\ a^4b^{-1}a^{-3}, a^5b^{-1}a^{-4}, a^6, a^5b\rangle$$

4.2 Practical computation

Let H be a finitely generated subgroup of the free group. By Corollary 4.2 $Cl(H)$ is the intersection of all the p-closures of H, and each of these is an overgroup of H (see the end of Sec. 2.1) by Corollary 2.18. So $Cl_{nil}(H)$ is also an overgroup, and it is the intersection of all the overgroups of H which are p-closed for some prime number p. Naturally, it suffices to take the intersection of the $\subseteq$-minimal overgroups of H which are p-closed for some p.

A procedure to compute the nil-closure of H is as follows:

(1) for each overgroup K of H, decide whether K is p-closed for some p; reject the overgroups K which are not p-closed for any p;

(2) compute the intersection of the overgroups K thus selected.

Since H has finitely many overgroups, we need only explain how to decide whether a given finitely generated subgroup of $F(A)$ is p-closed for some p.

As a first step, we prove the following.

Proposition 4.6 *Let H and K be finitely generated subgroups of the free group such that $H \subseteq K$. The set $Q(H, K)$ of all prime numbers p such that H is p-dense in K is empty or cofinite, and it is effectively computable.*

Proof. Since K is free, we first rewrite the generators of H as products of elements of a fixed basis of K. We are then reduced to the case where $K = F(A)$.

We slightly modify the procedure of Section 3.1. Let $\sigma\colon F(A) \to \mathbb{Z}^A$ be the natural morphism from the free group onto the free abelian group. Let

$h_1, \ldots, h_r$ be the given set of generators of H, and let $\mathfrak{M}(H)$ be the $r \times |A|$ matrix consisting of the row vectors $\sigma(h_1), \ldots, \sigma(h_r)$.

Next we compute the rank of $\mathfrak{M}(H)$ (in $\mathbb{Q}$). If that rank is not $|A|$, then every matrix $\mathfrak{M}_p(H)$ over $\mathbb{Z}/p\mathbb{Z}$ has rank less than $|A|$, so H is not p-dense for any p. In that case, $Q(H, F(A)) = \emptyset$.

If the rank of $\mathfrak{M}(H)$ is $|A|$, let d be the greatest common divisor of the non-zero order $|A|$ minor determinants of $\mathfrak{M}(H)$. For every prime p dividing d, $\mathfrak{M}_p(H)$ has rank less than $|A|$, so H is not p-dense. For every prime p not dividing d, one of the non-zero order $|A|$ minor determinants of $\mathfrak{M}(H)$ is not divisible by p, so $\mathfrak{M}_p(H)$ has rank $|A|$, and hence H is p-dense. Thus $Q(H, K)$ is the complement of the set of prime divisors of d, a cofinite set. $\square$

Our second step is to compute $\mathbb{P}(H)$, the set of prime numbers p such that H is p-closed. Let $\mathbb{P}$ be the set of all prime numbers.

Proposition 4.7 *Let H be a finitely generated subgroup of $F(A)$. The set $\mathbb{P}(H)$ is finite or cofinite, and it is effectively computable.*

Proof. We proceed by induction on the number of states of $\mathcal{A}(H)$. If $\mathcal{A}(H)$ has one state, then H is generated by a subset of the alphabet A, so H is p-closed for all p. That is, $\mathbb{P}(H) = \mathbb{P}$.

Let us now assume that $\mathcal{A}(H)$ has at least two states. By induction, we can compute $\mathbb{P}(K)$ for each non-trivial overgroup K of H, and each of these sets is finite or cofinite.

We define an equivalence relation on $\mathbb{P}$ by letting $p \equiv q$ whenever p and q belong to exactly the same sets $\mathbb{P}(K)$ (K overgroup properly containing of H). Then $\equiv$ has finite index, and its classes are the atoms of the (finite) Boolean algebra generated by the $\mathbb{P}(K)$. In particular, if a non-trivial overgroup K is p-closed, then it is q-closed for each q in the $\equiv$-class of p, $[p]_\equiv$.

If H admits $n \geq 2$ minimal (with respect to inclusion) p-closed non-trivial overgroups, say $K_1, \ldots, K_n$, then $Cl_p(H)$ is an overgroup of H contained in $K_1, \ldots, K_n$, so $Cl_p(H) = H$. If $q \equiv p$, the non-trivial overgroups of H are p-closed if and only if they are q-closed. So if H admits at least 2 minimal p-closed non-trivial overgroups, then H is q-closed for every q such that $q \equiv p$.

If H admits exactly one minimal p-closed non-trivial overgroup K, then for each $q \in [p]_\equiv$, the q-closure of H is either H or K, depending whether

H is q-dense in K. In that case, we have

$$[p]_\equiv \cap \mathbb{P}(H) = [p]_\equiv \setminus Q(H, K).$$

Thus $\mathbb{P}(H)$ is a finite Boolean combination of finite and cofinite sets, namely the $\mathbb{P}(K)$ and the $Q(H, K)$, which is effectively computable. □

Returning to the computation of the nil-closure of H, we need to consider the (minimal) overgroups K of H such that $\mathbb{P}(K) \neq \emptyset$. The nil-closure of H is the intersection of these overgroups:

Theorem 4.8 *Let H be a finitely generated subgroup of $F(A)$. The nil-closure of H is effectively computable.*

Note. The proof of Proposition 4.7 above gives an alternative method of computing the p-closure of H: consider the overgroups of H, starting from the greatest one, and for each decide whether it is p-closed. If H has several minimal non-trivial p-closed overgroups, then H is p-closed. If H has only one minimal non-trivial p-closed overgroup K, then $Cl_p(H)$ is K if H is p-dense in K, and H otherwise. This method may be interesting in practice, especially when $\mathcal{A}(H)$ has few, easily identifiable, quotients. We already noted a special case of this idea in Proposition 3.12 above.

In Examples 4.4 and 4.5, we were able to compute all the p-closures of H in a few operations, or more precisely to ascertain immediately that H was p-dense for almost all primes. This comes from the fact that the matrix $\mathfrak{M}(H)$ had rank $|A|$ (in $\mathbb{Q}$). This means that H was p-dense for all but a finite number of primes, namely the prime divisors of the gcd d of the non-zero order $|A|$ minor determinants of $\mathfrak{M}(H)$, and we needed only to calculate the p-closures of H for the prime divisors p of d. In the next example, the matrix $\mathfrak{M}(H)$ does not have rank $|A|$, and we use the procedure described above.

Example 4.9 Let us consider the pro-nilpotent closure of the subgroup

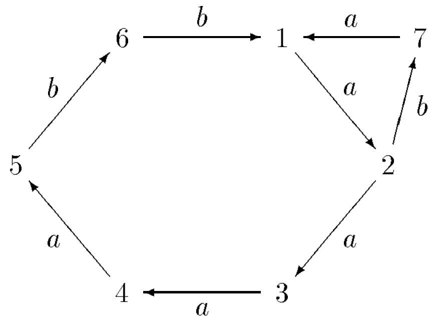

$H = \langle a^4b^2, aba \rangle$

Note that $\mathfrak{M}(H) = \begin{pmatrix} 4 & 2 \\ 2 & 1 \end{pmatrix}$: this matrix has rank 1, so H is not p-dense for any p. In order to determine all the p-closures of H, and hence the nil-closure of H, one needs to study the lattice of overgroups of H (dual of the lattice of congruences of $\mathcal{A}(H)$). One verifies that all 11 overgroups of H properly contained in $F(A)$ have rank bigger than H. By Proposition 3.12 and in view of the value of $\mathfrak{M}(H)$, it follows that H is p-closed for every p, and in particular H is nil-closed.

This example illustrates a more general situation where the computation is simpler.

Proposition 4.10 *Let H be a finitely generated subgroup of the free group $F(A)$. If an order $|A|$ minor of $\mathfrak{M}(H)$ has determinant ± 1, then H is p-dense for every p, and hence H is nil-dense. If no order $|A|$ minor of $\mathfrak{M}(H)$ has determinant ± 1 and every overgroup of H properly contained in $F(A)$ has rank greater than H, then H is p-closed for some p and H is nil-closed.*

5 Decidability results for pseudovarieties of inverse monoids

We now apply the results of the previous sections to the theory of inverse monoids. A *pseudovariety of inverse monoids* is a class $\mathbf{V}$ of finite inverse monoids which is closed under taking inverse submonoids, homomorphic images and finite direct products. For instance, the class $\mathbf{SL}$ of commutative idempotent monoids is such a pseudovariety. Every pseudovariety of groups can also be considered to be a pseudovariety of inverse monoids. If $\mathbf{V}$ is a pseudovariety of inverse monoids, we let $\mathbf{SL} \circledm \mathbf{V}$ be the class of finite inverse monoids M such that there exist inverse monoids R and V and morphisms

$\alpha: R \to M$ and $\beta: R \to V$ with α onto, $V \in \mathbf{V}$ and $\alpha\beta^{-1}(1) \in \mathbf{SL}$. The class $\mathbf{SL} \textcircled{m} \mathbf{V}$, called the *Mal'cev product of* $\mathbf{SL}$ *and* $\mathbf{V}$, is easily seen to be a pseudovariety of inverse monoids. In this section, we will prove the decidability of the membership problem for certain pseudovarieties of the form $\mathbf{SL} \textcircled{m} \mathbf{V}$.

Before we can prove this result, we need to introduce some more machinery, and connect the notion of $\mathbf{V}$-extendible subgroups of the free group with the pro-identities satisfied by $\mathbf{V}$.

5.1 Extendible subgroups and pro-identities of V

We have seen in Section 2.1 that there exists a free inverse monoid over each alphabet A, denoted $FIM(A)$. Let $\rho: (A \cup A^{-1})^* \to FIM(A)$ be the natural onto morphism. As in Section 1.1, we can consider the projective limit of all A-generated finite inverse monoids. It is the profinite completion $\widehat{FIM}(A)$ of $FIM(A)$, and it is the free profinite inverse monoid over A. Then $FIM(A)$ is dense in that compact monoid and we let $\jmath: FIM(A) \to \widehat{FIM}(A)$ be the (one-to-one) inclusion morphism (see [2]).

By construction of $\widehat{FIM}(A)$, every morphism from $FIM(A)$ into a finite inverse monoid extends uniquely to a continuous morphism defined on $\widehat{FIM}(A)$. If $\mathcal{A}$ is an inverse automaton with transition morphism $\mu: (A \cup A^{-1})^* \to M(\mathcal{A})$, then $M(\mathcal{A})$ is a finite inverse monoid, so μ induces a morphism $FIM(A) \to M(\mathcal{A})$, and a continuous morphism $\widehat{FIM}(A) \to M(\mathcal{A})$, which are also denoted by μ. In particular, it makes sense to talk of transitions of $\mathcal{A}$ induced by elements of $\widehat{FIM}(A)$ and to write $q \cdot u$ if q is a state of $\mathcal{A}$ and $u \in \widehat{FIM}(A)$.

Since each (A-generated) group is an (A-generated) inverse monoid, there is a natural onto morphism $\gamma: FIM(A) \to F(A)$, and if $\mathbf{V}$ is a pseudovariety of groups, there is a natural continuous onto morphism σ, from $\widehat{FIM}(A)$ onto $\hat{F}_A(\mathbf{V})$. Then every continuous morphism from $\widehat{FIM}(A)$ into a group of $\mathbf{V}$ can be factored through σ.

Let $u \in \widehat{FIM}(A)$. We say that a finite group G *satisfies the pseudoidentity* $u = 1$ if for each continuous morphism $\varphi: \widehat{FIM}(A) \to G$, we have $\varphi(u) = 1$. We say that the pseudovariety $\mathbf{V}$ satisfies $u = 1$ if all the elements of $\mathbf{V}$ satisfy $u = 1$. This is equivalent to the equality $\sigma(u) = 1$.

Let H be a finitely generated subgroup of $F(A)$. We define the relation $\sim_{\mathbf{V}}$ on the set of states of $\mathcal{A}(H)$ by letting $p \sim_{\mathbf{V}} q$ if and only if there exists $u \in \widehat{FIM}(A)$ such that $p \cdot u = q$ and $\mathbf{V}$ satisfies $u = 1$.

Proposition 5.1 *Let H be a finitely generated subgroup of $F(A)$. The congruences $\sim_{\mathbf{V}}$ and $\sim$ on $\mathcal{A}(H)$ coincide.*

Proof. Let us fix, for each state p of $\mathcal{A}(H)$, a reduced word u_p such that $1 \cdot u_p = p$ in $\mathcal{A}(H)$.

Let us first assume that $p \sim_{\mathbf{V}} q$. Let K be a clopen subgroup of $F(A)$ containing H such that $\sim = \sim_{H,K}$. In particular, $\mathcal{A}(\tilde{H}) = \mathcal{A}(H)/\sim$ embeds in $\mathcal{A}(K)$. By hypothesis, there exists an element $u \in \widehat{FIM}(A)$ such that $p \cdot u = q$ and $\mathbf{V}$ satisfies $u = 1$. In particular, we have $1 \cdot (u_p u u_q^{-1}) = 1$ in $\mathcal{A}(H)$, so the same equality holds in $\mathcal{A}(\tilde{H})$, and hence also in $\mathcal{A}(K)$.

Moreover, since K is clopen, we have $M(K) \in \mathbf{V}$, so $M(K)$ satisfies $u = 1$, that is, $r \cdot u = r$ for each state r of $\mathcal{A}(K)$. Therefore $1 \cdot u_p u_q^{-1} = 1$ in $\mathcal{A}(K)$, that is, $p \sim_{H,K} q$, and hence $p \sim q$.

To prove the converse, let us now assume that $p \sim q$. We have the following commutative diagram, where φ is an arbitrary morphism from $F(A)$ into a group $V \in \mathbf{V}$ and $\hat{\varphi}$ is the continuous morphism from $\hat{F}_A(\mathbf{V})$ into V induced by φ.

$$\begin{tikzcd}
 & \widehat{FIM}(A) \arrow[r, "\sigma"] & \hat{F}_A(\mathbf{V}) \arrow[dr, "\hat{\varphi}"] & \\
(A \cup A^{-1})^* \arrow[r, "\rho"'] & FIM(A) \arrow[u, "\jmath"] \arrow[r, "\gamma"'] & F(A) \arrow[u, "\imath"] \arrow[r, "\varphi"'] & V
\end{tikzcd}$$

Since $p \sim q$, we have $u_p u_q^{-1} \in Cl(H)$, or more precisely, $\gamma\rho(u_p u_q^{-1}) \in Cl(H)$. But $Cl(H) = \imath^{-1}(\overline{\imath(H)})$ (see Section 1.1), so $\imath\gamma\rho(u_p u_q^{-1}) \in \overline{\imath(H)}$. A basis of neighborhoods of $\imath\gamma\rho(u_p u_q^{-1})$ in $\hat{F}_A(\mathbf{V})$ is given by the sets of the form $\hat{\varphi}^{-1}\hat{\varphi}\imath\gamma\rho(u_p u_q^{-1})$, where φ runs over the morphisms from $F(A)$ into groups $V \in \mathbf{V}$. Thus, for each such morphism φ, there exists a word $h_\varphi \in (A \cup A^{-1})^*$ such that $\gamma\rho(h_\varphi) \in H$ and $\hat{\varphi}\imath\gamma\rho(h_\varphi) = \hat{\varphi}\imath\gamma\rho(u_p u_q^{-1})$. Then $\hat{\varphi}\imath\gamma\rho(u_p^{-1} h_\varphi u_q) = 1$, that is, $\hat{\varphi}\sigma\jmath\rho(u_p^{-1} h_\varphi u_q) = 1$.

Let h be a limit point of the $\jmath\rho(h_\varphi)$ in $\widehat{FIM}(A)$ (the existence of such a limit point is ensured by compactness). Then $\sigma\jmath\rho(u_p^{-1})\sigma(h)\sigma\jmath\rho(u_q) = 1$, and hence $\mathbf{V}$ satisfies $\jmath\rho(u_p^{-1})h\jmath\rho(u_q) = 1$.

In addition, for each morphism $\varphi: F(A) \to V$ $(V \in \mathbf{V})$, $\gamma\rho(h_\varphi) \in H$, so $1 \cdot \jmath\rho(h_\varphi) = 1$ in $\mathcal{A}(H)$, and hence $1 \cdot h = 1$. Therefore, $p \cdot \jmath\rho(u_p^{-1})h\jmath\rho(u_q) = q$, that is, $p \sim_{\mathbf{V}} q$. $\square$

5.2 The memberhip problem for $\mathbf{SL} ⓜ \mathbf{V}$ when $\mathbf{V}$ is a pseudovariety of groups

We recall the following result, a special case of [14].

Proposition 5.2 *Let $\mu: \widehat{FIM}(A) \to M$ be a continuous morphism onto a finite inverse monoid M. Let $\mathbf{V}$ be a pseudovariety of groups and let $\sigma: \widehat{FIM}(A) \to \hat{F}_A(\mathbf{V})$ be the natural onto continuous morphism. Then $M \in \mathbf{SL} ⓜ \mathbf{V}$ if and only if $\mu\sigma^{-1}(1) \in \mathbf{SL}$.*

Then we have the following result.

Proposition 5.3 *Let H be a finitely generated subgroup of the free group $F(A)$ and let $\mathbf{V}$ be a pseudovariety of groups. Then $M(H) \in \mathbf{SL} ⓜ \mathbf{V}$ if and only if H is $\mathbf{V}$-extendible.*

Proof. Let μ be the continuous extension of the transition morphism of $\mathcal{A}(H)$, $\mu: \widehat{FIM}(A) \to M(H)$, and let σ be the natural morphism $\sigma: \to \hat{F}_A(\mathbf{V})$. By Proposition 5.2, $M(H) \in \mathbf{SL} ⓜ \mathbf{V}$ if and only if $\mu\sigma^{-1}(1) \in \mathbf{SL}$. Recall that in an inverse monoid, the idempotents always commute. So $M(H) \in \mathbf{SL} ⓜ \mathbf{V}$ if and only if, for each $u \in \widehat{FIM}(A)$, $\sigma(u) = 1$ implies $\mu(u) = \mu(u^2)$.

But $\sigma(u) = 1$ means that $\mathbf{V}$ satisfies the pro-identity $u = 1$ (see Section 5.1), so $M(H) \in \mathbf{SL} ⓜ \mathbf{V}$ if and only if whenever $\mathbf{V}$ satisfies $u = 1$ ($u \in \widehat{FIM}(A)$), then $\mu(u)$ is idempotent. Now to say that $\mu(u)$ is idempotent means that for each state q of $\mathcal{A}(H)$ such that $q \cdot u$ exists, we have $q \cdot u^2 = q \cdot u$, and hence, since $\mathcal{A}(H)$ is an inverse automaton, $q \cdot u = q$ each time that $q \cdot u$ is defined.

Thus $M(H) \in \mathbf{SL} ⓜ \mathbf{V}$ if and only if whenever $\mathbf{V}$ satisfies $u = 1$, then for each state q of $\mathcal{A}(H)$ such that $q \cdot u$ is defined, we have $q \cdot u = q$. This states exactly that the automaton congruence $\sim_{\mathbf{V}}$ (defined in Section 5.1) is trivial. By Proposition 5.1 it follows that $M(H) \in \mathbf{SL} ⓜ \mathbf{V}$ if and only if the congruence $\sim$ is trivial on $\mathcal{A}(H)$, that is, if and only if H is $\mathbf{V}$-extendible. $\square$

Corollary 5.4 *If $\mathbf{V}$ is extension closed, then $M(H) \in \mathbf{SL} ⓜ \mathbf{V}$ if and only if H is closed.*

Proof. Immediate by Propositions 5.3 and 2.17. $\square$

We now turn to the membership problem in $\mathbf{SL} ⓜ \mathbf{V}$.

Proposition 5.5 *Let* $\mathbf{V}$ *be a pseudovariety of groups. If one can decide whether a finitely generated subgroup of the free group is* $\mathbf{V}$*-extendible, then the membership problem in* $\mathbf{SL} \textcircled{m} \mathbf{V}$ *is decidable.*

If, given a finitely generated subgroup H *of the free group* $F(A)$*, the membership problem of the* $\mathbf{V}$*-closure of* H *is decidable, then the membership problem of* $\mathbf{SL} \textcircled{m} \mathbf{V}$ *is decidable.*

Proof. It is known that every finite inverse monoid is a subdirect product of transition monoids of finite inverse automata. Observe that not every finite inverse automaton is reduced in the sense of this paper, since it may have states of degree one besides the distinguished state. However, we now show that if $\mathcal{A}$ is any finite inverse automaton, then we can find a reduced finite inverse automaton $\mathcal{B}$ such that the transition monoid of $\mathcal{A}$ is isomorphic to that of $\mathcal{B}$.

Let $\mathcal{A}$ be a finite inverse automaton over the alphabet A. Now consider the set $B = A \cup A^{-1}$ as an alphabet. Let $\mathcal{B}$ be the inverse automaton over the alphabet B with the same state set and the same initial-terminal state as $\mathcal{A}$, and such that $q \cdot y$ has the same value as in $\mathcal{A}$, for each letter $y \in B = A \cup A^{-1}$. Now in $\mathcal{B}$, all states have degree at least 2, so $\mathcal{B}$ is reduced. Moreover it is immediate that $\mathcal{A}$ and $\mathcal{B}$ have the same transition monoids so we are done.

It follows now that every finite inverse monoid is a subdirect product of transition monoids of the form $M(H)$, where H is a finitely generated subgroup of the free group. So the first statement follows immediately from Proposition 5.3. The second statement is a consequence of Proposition 3.3. $\square$

Corollary 5.6 *The membership problems for* $\mathbf{SL} \textcircled{m} \mathbf{G}_p$ *(p a prime number) and* $\mathbf{SL} \textcircled{m} \mathbf{G}_{nil}$ *are decidable.*

Conclusion

We have separated the concepts of $\mathbf{V}$-extendibility and of $\mathbf{V}$-closure for finitely generated subgroups of the free groups. These notions are equivalent when $\mathbf{V}$ is an extension-closed pseudovariety, an hypothesis made in all previous works on the computation of $\mathbf{V}$-closures.

In addition, we have simplified and sped up Ribes and Zalesskiĭ's algorithm for the computation of the pro-p closure of a finitely generated subgroup of the free group, showing that this computation can be performed in polynomial time.

Next we have extended this result to prove the computability of the pro-nil closure of a finitely generated subgroup of the free group. Finally, we have considered an application of our results to a problem in finite monoid theory, the membership problem in pseudovarieties of inverse monoids which are Mal'cev products of semilattices and a pseudovariety of groups $\mathbf{V}$. Other applications and connections with monoid and language theory are discussed in the introduction.

We conclude this paper by pointing to our readers the case of solvable groups. The class of finite solvable groups is an extension-closed pseudovariety, and hence general results apply to it: the pro-solvable closure of a finitely generated subgroup H of the free group is an overgroup of H, which has rank at most equal to the rank of H. Thus it is one of a finite computable list of subgroups, but we do not know how to determine which overgroup of H is its pro-solvable closure. In fact, it would suffice to have an algorithm to decide whether a given finitely generated subgroup of the free group is dense in the pro-solvable topology: if we were able to decide this, then we would be able to compute the pro-solvable closure of H as we did for the p-closure in the remark following Theorem 4.8. At the moment, this problem is wide open.

References

[1] J. Almeida. *Finite semigroups and universal algebra*, World Scientific Publishing, Singapore, 1994.

[2] J. Almeida and P. Weil. Relatively free profinite monoids, in J. Fountain ed. *Semigroups, formal languages and groups*, NATO ASI Series **C-466**, Kluwer, Dordrecht, 1995.

[3] C.J. Ash. Inevitable graphs: a proof of the type II conjecture and some related decision procedures, *Int. J. Algebra and Comput.* **1** (1991) 411–436.

[4] N. Bourbaki. *Eléments de mathématiques, topologie générale*, Hermann, Paris, 1971.

[5] J.D. Dixon, M.P.F. du Sautoy, A. Mann and D. Segal. *Analytic pro-p groups*, London Mathematical Society Lecture Notes **157**, Cambridge University Press, 1991.

[6] M. Fried and M. Jarden. *Field arithmetic*, Springer, Berlin, 1986.

[7] D. Gildenhuys and L. Ribes. A Kurosh subgroup theorem for free pro-$\mathcal{C}$ products of pro-$\mathcal{C}$ groups, *Trans. Amer. Math. Soc.* **186** (1973) 309–329.

[8] M. Hall Jr. Coset representation in free groups, *Trans. Amer. Math. Soc.* **67** (1949) 421–432.

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

[10] K. Henckell, S.W. Margolis, J.-E. Pin and J. Rhodes. Ash's type II theorem, profinite topology and Malcev products, *Int. J. Algebra and Comput.* **1** (1991) 411–436.

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

[12] M. Petrich. *Inverse semigroups*, Wiley, New York, 1984.

[13] J.-E. Pin. Topologies for the free monoid, *J. Algebra* **137** (1991) 297–337.

[14] J.-E. Pin and P. Weil. Profinite semigroups, Mal'cev products and identities, *J. Algebra* **182** (1996) 604–626.

[15] F. Rakoczi. Fast recognition of the nilpotency of permutation groups, in *Proc. International Symposium on Symbolic and Algebraic Computation (ISSAC)*, Montréal, 1995, 265–269.

[16] L. Ribes. *Introduction to profinite groups and Galois cohomology*, Queen's papers in pure and applied mathematics **24**, Queen's university, Kingston, 1970.

[17] L. Ribes and P.A. Zalesskiĭ. On the profinite topology on a free group, *Bull. London Math. Soc.* **25** (1993) 37–43.

[18] 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.

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

[20] J.-P. Serre. *Cohomologie galoisienne*, Lecture Notes in mathematics **5**, Springer, 1965.

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

[22] B. Steinberg. Inevitable graphs and profinite topologies: some solutions to algorithmic problems in monoid and automata theory, stemming from group theory, submitted.