

THE GEOMETRY OF PROFINITE GRAPHS WITH APPLICATIONS TO FREE GROUPS AND FINITE MONOIDS

K. AUINGER AND B. STEINBERG

ABSTRACT. We initiate the study of the class of profinite graphs Γ defined by the following geometric property: for any two vertices v and w of Γ , there is a (unique) smallest connected profinite subgraph of Γ containing them; such graphs are called *tree-like*. Profinite trees in the sense of Gildenhuys and Ribes are tree-like, but the converse is not true. A profinite group is then said to be *dendral* if it has a tree-like Cayley graph with respect to some generating set; a Bass-Serre type characterization of dendral groups is provided. Also, such groups (including free profinite groups) are shown to enjoy a certain small cancellation condition.

We define a pseudovariety of groups \mathbf{H} to be *arboreous* if all finitely generated free pro- \mathbf{H} groups are dendral (with respect to a free generating set). Our motivation for studying such pseudovarieties of groups is to answer several open questions in the theory of profinite topologies and the theory of finite monoids. We prove, for arboreous pseudovarieties \mathbf{H} , a pro- \mathbf{H} analog of the Ribes and Zalesskiĭ product theorem for the profinite topology on a free group. Also, arboreous pseudovarieties are characterized as precisely the solutions \mathbf{H} to the much studied pseudovariety equation $\mathbf{J} * \mathbf{H} = \mathbf{J} \textcircled{\cap} \mathbf{H}$.

1. INTRODUCTION

In the early nineties, the work surrounding the Rhodes Type II conjecture [20] led to an unexpected interplay between two seemingly disjoint fields: profinite group theory and finite monoid theory. Pin and Reutenauer [36] proved that the truth of this conjecture would be implied by the validity of the following conjecture on the profinite topology of a free group F : if H_1, \dots, H_n are finitely generated subgroups of F , then the product $H_1 \cdots H_n$ is closed in F with respect to the profinite topology. The case $n = 1$ is Hall's celebrated theorem [18].

In the end, the Type II conjecture was proved independently (and at virtually the same time) by Ash [8] and Ribes and Zalesskiĭ [39]. Ash's proof used inverse monoids, automata, and geometric arguments involving Cayley graphs of finite groups to obtain the result directly. Ribes and Zalesskiĭ proved the conjecture

Received by the editors September 28, 2002 and, in revised form, March 13, 2003.

2000 *Mathematics Subject Classification*. Primary 20E18, 20E08, 20M07, 20M18.

The authors gratefully acknowledge support from INTAS project 99-1224 *Combinatorial and geometric theory of groups and semigroups and its applications to computer science*. The second author was supported in part by NSF-NATO postdoctoral fellowship DGE-9972697, as well as by Praxis XXI scholarship BPD 16306 98, FCT through the *Centro de Matemática da Universidade do Porto*, and by the FCT and POCTI approved project POCTI/32817/MAT/2000 in participation with the European Community Fund FEDER. He was at the University of Porto when this work was performed.

of Pin and Reutenauer using the theory of profinite groups acting on profinite graphs [39]. They even were able to prove the following more general result [40]: if \mathbf{H} is a *pseudovariety* of groups (that is, a class of finite groups closed under taking finite direct products, subgroups and quotients) which, in addition, is closed under extensions and H_1, \dots, H_n are finitely generated subgroups of a free group F which are closed in the pro- \mathbf{H} topology of F , then the product $H_1 \cdots H_n$ is also closed. We say that a pseudovariety \mathbf{H} of groups satisfying this latter property (for n subgroups) enjoys *n -coset separability* (motivated by the terminology *double coset separability* [16]).

One key property used by Ribes and Zalesskiĭ was that the Cayley graphs of the finitely generated free pro- \mathbf{H} groups are \mathbf{H} -trees (where such are defined by a profinite version of the usual homological characterization of a tree) when \mathbf{H} is extension-closed. This led Almeida and Weil [5], [6] to define a pseudovariety of groups \mathbf{H} to be *arborescent* if all finitely generated free pro- \mathbf{H} groups have Cayley graphs which are \mathbf{H} -trees. They showed that this happens if and only if \mathbf{H} is closed under co-extension by its Abelian members. They also obtained some factorization results that allowed them to solve various problems in monoid theory.

At this point in time, the interplay between finite monoid theory and profinite group theory was in one direction only. Then in [52] the second author showed, via a translation into inverse monoid theory, that a key intermediary result of Ash in his proof of the Type II conjecture was formally equivalent to the product theorem of Ribes and Zalesskiĭ. What is more, using this translation he could show that arborescent pseudovarieties enjoy double coset separability. We should mention that Ribes and Zalesskiĭ's proof [39], [40] only applies in the extension-closed setting, as they used the property (which only holds for such pseudovarieties) that the subspace topology on an open subgroup H of a free group F (endowed with the pro- \mathbf{H} topology) coincides with its own (full) pro- \mathbf{H} topology. So the result was truly novel and was a first indication that monoid theory had something to say about group theory.

Starting where [52] left off, this paper uses a *mélange* of profinite group theory and inverse monoid theory to prove a generalization of Ribes and Zalesskiĭ's product theorem for the pro- \mathbf{H} topology on the free group to a much larger class of pseudovarieties \mathbf{H} . In the process, we obtain a new proof of that theorem for the case $\mathbf{H} = \mathbf{G}$, the pseudovariety of all finite groups.

The principal property of \mathbf{H} -trees used in the various results above is that any two vertices are contained in a unique minimal connected subgraph. We thus define a profinite graph to be *tree-like* if it has this property. The first part of the paper explores the basic geometric properties of tree-like graphs. Let us call a profinite group *dendral with respect to a (profinite) set of generators* if it has a tree-like Cayley graph with respect to that generating set. More generally, we shall say that a profinite group is *dendral* if it is dendral with respect to some generating set. For example, finitely generated free profinite groups are dendral. We give a Bass-Serre type characterization of such groups in terms of free actions on tree-like graphs. In particular, open subgroups of dendral groups are shown to be dendral. Next we prove a cancellation result for dendral groups which is a direct analog to the following obvious property of a free group: if w_1, \dots, w_n are elements of a free group F , $n \geq 3$, and if, for all $i \in \{2, \dots, n-1\}$, some letter of w_i remains uncanceled in the reduced form of $w_{i-1}w_iw_{i+1}$, then $w_1 \cdots w_n \neq 1$ in F .

We call a pseudovariety of groups \mathbf{H} *arboreous* if all finitely generated free pro- \mathbf{H} groups are dendral (with respect to a free generating set). Non-trivial arboreous pseudovarieties residually contain the free group and are join irreducible. Moreover, the union of an upwardly directed set of arboreous pseudovarieties is again arboreous; in particular, this holds for the union of an ascending chain. Arboreous pseudovarieties are, of course, arboreous — the converse is not true.

To investigate the nature of arboreous pseudovarieties, we use techniques from inverse monoid theory. Recall that $\mathbf{SI} * \mathbf{H}$ is the pseudovariety of all finite inverse monoids with E -unitary covers over \mathbf{H} . We generalize the results of [48], [52] by showing that \mathbf{H} is arboreous if and only if the free pro- $\mathbf{SI} * \mathbf{H}$ inverse monoids on finite sets are F -inverse. It is also proved that \mathbf{H} is arboreous if and only if the \mathbf{H} -pointlike sets of a finite inverse monoid in $\mathbf{SI} * \mathbf{H}$ are the \mathbf{G} -pointlike sets.

In [52], a pseudovariety of groups was called *Hall* if each \mathbf{H} -extendible finitely generated subgroup of a free group in the sense of [29] is pro- \mathbf{H} closed. Equivalently, \mathbf{H} is Hall if and only if the fundamental group of each connected subgraph of a Cayley graph of a group in \mathbf{H} is closed in the appropriate free group with respect to the pro- \mathbf{H} topology. In the main result of the paper, we prove that a pseudovariety \mathbf{H} is arboreous if and only if it is Hall and enjoys n -coset separability for all n , if and only if it is Hall and satisfies double-coset separability. In particular, each arboreous pseudovariety \mathbf{H} enjoys n -coset separability for each n . The proof uses the aforementioned cancellation result and the techniques of [52]; however, its primary inspiration comes from Ash's proof of the Type II conjecture, and it can be viewed as transferring some ideas of Ash's paper [8] to the profinite setting.

We then follow with a section that proves that the set of non-trivial arboreous pseudovarieties forms a right ideal in the monoid of all pseudovarieties of groups — the binary operation on the set of group pseudovarieties here is the Mal'cev (or semidirect) product. It follows that the set of arboreous pseudovarieties is much larger than the set of arboreous pseudovarieties.

Using a co-extension of the sort developed in [15], [38], [31], we are able to show that a sufficient condition for a pseudovariety \mathbf{H} to be arboreous is the following: for each $G \in \mathbf{H}$, there is a cyclic group C such that the wreath product $C \wr G$ is in \mathbf{H} . This is a reasonably weak condition, and we exhibit an example of such a pseudovariety, all of whose cyclic members have square-free order. In fact the pseudovariety of all groups with square-free exponent satisfies this condition.

The Type II conjecture was concerned with one of the fundamental areas of research in finite monoid theory: the study of co-extensions. For instance, if one has an exact sequence

$$1 \rightarrow N \rightarrow G \rightarrow H \rightarrow 1$$

of groups, then there is the notion of the sequence being split. Essentially, the sequence splits if G is a semidirect product of N by H . A classical result of group theory [33] asserts that, for any exact sequence (as above), G embeds in the wreath product $N \wr H$, which in turn is a semidirect product of the form $N^H \rtimes H$. Thus, at the pseudovariety level, split co-extensions and co-extensions of groups are the same thing. This is not the case in monoid theory.

The pseudovariety generated by all co-extensions of groups in \mathbf{H} by monoids in \mathbf{V} is the Mal'cev product $\mathbf{V} \circledast \mathbf{H}$ (for a precise definition of co-extension see section 8). This contains the pseudovariety $\mathbf{V} * \mathbf{H}$ generated by all semidirect products of members of \mathbf{V} by members of \mathbf{H} , but in general these pseudovarieties are not the

same. A sufficient condition for equality is that \mathbf{V} be local in the sense of Tilson [56]; this happens, for instance, if \mathbf{V} is the pseudovariety \mathbf{Sl} of semilattices. The Type II conjecture proposed an algorithm to decide membership in pseudovarieties of the form $\mathbf{V} \circledast \mathbf{G}$ for pseudovarieties \mathbf{V} with decidable membership problem. Interest in semidirect products of pseudovarieties arose out of the Krohn-Rhodes theorem [23] and the theory of complexity of finite semigroups [24], although the formulation of these notions in terms of semidirect products of pseudovarieties is due to Eilenberg [14].

The much studied pseudovariety \mathbf{J} of \mathcal{J} -trivial monoids is not local. However, a famous result of Henckell and Rhodes [20], using Ash's theorem [8], shows that $\mathbf{J} * \mathbf{G} = \mathbf{J} \circledast \mathbf{G}$. Pin asked in [35] whether such is true for the pseudovarieties of p -groups and solvable groups. This led the second author to undertake in [50] a comprehensive study of the pseudovariety equation

$$(1.1) \quad \mathbf{J} * \mathbf{H} = \mathbf{J} \circledast \mathbf{H}$$

for pseudovarieties of groups \mathbf{H} . In particular, equality was shown to hold if \mathbf{H} is closed under extension; this was improved in [52] to include all arborescent pseudovarieties. Also it was shown in [50] that equality can fail if \mathbf{H} is a non-trivial pseudovariety of Abelian groups or of bounded exponent.

The pseudovariety $\mathbf{J} * \mathbf{H}$ is of special importance since it corresponds, in formal language theory, to the Boolean polynomial closure of the languages recognized by \mathbf{H} [51]. Since, in general, membership in $\mathbf{J} \circledast \mathbf{H}$ is easier to compute than membership in $\mathbf{J} * \mathbf{H}$, equality is quite useful to have from a computational standpoint. On the other hand, semidirect products are easier to understand algebraically in terms of their factors than more general co-extensions, so equality is also useful for understanding $\mathbf{J} \circledast \mathbf{H}$ from a structural point of view. In section 8 we show that (1.1) holds precisely for the class of arboreous pseudovarieties of groups, thereby completely settling the question of when all \mathbf{J} -co-extensions of a pseudovariety of groups “split.”

Other applications include the following. Ever since Ash proved the Type II conjecture, J.-E. Pin has been asking at various conferences for a direct construction that, given a member M of $\mathbf{J} \circledast \mathbf{G}$, finds a finite group G , whose order and structure are tightly determined by M , such that M divides a semidirect product of the form $J * G$ with $J \in \mathbf{J}$. We provide such a construction which gives a group G that is “small,” both in terms of the size $|G|$ of G , and the pseudovariety $\langle G \rangle$ generated by G .

The paper ends with an appendix characterizing which pseudovarieties of groups \mathbf{H} have the property that there is a profinite ring R with unity such that the Cayley graphs of the finitely generated free pro- \mathbf{H} groups are R -trees. This occurs if and only if there is a prime p such that \mathbf{H} is closed under co-extensions by p -groups. In particular, there are arboreous pseudovarieties which do not enjoy this property.

All groups, monoids and semigroups occurring in this paper, except (relatively) free and profinite ones, are assumed to be finite, unless otherwise stated.

2. PROFINITE GRAPHS

We follow the Serre [44] convention and define a *graph* Γ to consist of a set $V(\Gamma)$ of *vertices* and disjoint sets $E(\Gamma)$ of *positively oriented (or positive) edges* and $E^{-1}(\Gamma)$ of *negatively oriented (or negative) edges* together with *incidence* functions

$\iota, \tau : E(\Gamma) \cup E(\Gamma)^{-1} \rightarrow V(\Gamma)$ selecting the *initial*, respectively, *terminal* vertices of an edge e and mutually inverse bijections (both written: $e \mapsto e^{-1}$) between $E(\Gamma)$ and $E^{-1}(\Gamma)$ such that $\iota e^{-1} = \tau e$ for all edges e (whence $\tau e^{-1} = \iota e$, as well). We set $\widetilde{E}(\Gamma) = E(\Gamma) \cup E^{-1}(\Gamma)$ and call it the *edge set* of Γ . Given this definition of a graph, the notions of *subgraph*, *direct product of graphs*, *morphism of graphs* and *projective limit of graphs* have the obvious meanings. Edges are to be thought of geometrically: when one draws an oriented graph, one draws only the edge e and thinks of e^{-1} as being the same edge, but traversed in the reverse direction.

A *path* p in a graph Γ is a finite sequence $p = e_1 \dots e_n$ of consecutive edges, that is, $\tau e_i = \iota e_{i+1}$ for all i ; we define $\iota p = \iota e_1$ to be the initial vertex of p and $\tau p = \tau e_n$ to be the terminal vertex of p . The path $p^{-1} = e_n^{-1} \dots e_1^{-1}$ is the *inverse* of the path p . We also consider an empty path at each vertex. Two paths p and q are *co-terminal* if they have the same initial and terminal vertices, respectively. A path $p = e_1 \dots e_n$ is a *circuit* with base point v if $v = \iota p = \tau p$. A circuit $p = e$ consisting of one edge only is a *loop*. A graph is *connected* if any two vertices can be joined by a path. In particular, the empty graph is connected.

If $V(\Gamma)$ and $E(\Gamma)$ are topological spaces, $\widetilde{E}(\Gamma)$ is the topological sum of $E(\Gamma)$ and $E^{-1}(\Gamma)$, and ι, τ and $(\)^{-1}$ are continuous, then Γ is called a *topological graph*. A *profinite graph* is a topological graph Γ which is a projective limit of finite, discrete graphs. It is well known (see [59], [41]) that Γ is profinite if and only if $V(\Gamma)$ and $\widetilde{E}(\Gamma)$ are both compact, totally disconnected Hausdorff spaces (sometimes called Boolean spaces). We shall usually refer to such spaces as *profinite sets*. Morphisms among profinite graphs are always assumed to be continuous. All compact spaces occurring in the paper shall be assumed Hausdorff.

The profinite graphs of primary interest are the Cayley graphs of relatively free profinite groups, where a *profinite group* is a compact, totally disconnected group, or, equivalently, a projective limit of finite groups. We refer the reader to [42], [43], [6] for basic definitions on profinite and relatively free profinite groups.

For an alphabet A , we use A^{-1} to denote a disjoint copy of A , and \widetilde{A} for the disjoint union of A and A^{-1} ; that is, $\widetilde{A} = A \cup A^{-1}$. If, in addition, A is a topological space, then A^{-1} is a disjoint homeomorphic copy of A , and \widetilde{A} is the topological sum of A and A^{-1} . Suppose that G is a topological group, topologically generated by a topological space A ; that is, there is a continuous map $\psi : A \rightarrow G$ such that $\overline{\langle A\psi \rangle} = G$. Then the *Cayley graph* of G with respect to A , denoted by $\Gamma_A(G)$, has vertex set G , edge set $G \times \widetilde{A}$, incidence functions given by $\iota(g, a) = g$, $\tau(g, a) = g \cdot a\psi$ and involution $(g, a)^{-1} = (g \cdot a\psi, a^{-1})$. We call $a \in \widetilde{A}$ the *label* of (g, a) . The topology on the vertex set is the topology of G ; the edge set receives the product topology of $G \times \widetilde{A}$. If G is a profinite group and A is a profinite set, then $\Gamma_A(G)$ is a profinite graph. Note that G has a natural, continuous left action on $\Gamma_A(G)$. Moreover, if A is (pro)finite and G is A -generated profinite with $G = \varprojlim G_i$, the groups G_i forming an inverse system of finite quotients of G , then $\Gamma_A(G) = \varprojlim \Gamma_A(G_i)$. From now on, when we speak of generators of a topological algebra, we shall always mean topological generators.

A *pseudovariety* of algebras \mathbf{V} is a class of finite algebras closed under the operations of forming finite direct products, taking subalgebras, and taking morphic images [1], [14]. For the moment, we shall be concerned with pseudovarieties of groups;

later we shall turn towards pseudovarieties of inverse monoids and of monoids. Examples include the pseudovarieties: \mathbf{G} of all finite groups, \mathbf{G}_p of all finite p -groups (p a prime), \mathbf{Ab} of all finite Abelian groups, \mathbf{G}_{nil} of all finite nilpotent groups, \mathbf{G}_{sol} of all finite solvable groups, and \mathbf{Ab}_n of all finite Abelian groups of exponent dividing n .

If \mathbf{H} is a pseudovariety of groups, then a profinite group G is said to be *pro- \mathbf{H}* if it is a projective limit of groups in \mathbf{H} ; equivalently, its finite (continuous) images all belong to \mathbf{H} . For each pro-finite set A there is a (relatively) free A -generated pro- \mathbf{H} group [42]. We shall denote this group $\widehat{F}_{\mathbf{H}}(A)$. It enjoys the following universal property: given any pro- \mathbf{H} group G and any continuous map $\psi : A \rightarrow G$, there is a unique continuous morphism $\widehat{\psi} : \widehat{F}_{\mathbf{H}}(A) \rightarrow G$ extending ψ . The notation is motivated as follows. Let A be finite, discrete and let $F_{\mathbf{H}}(A)$ be the relatively free group on A in the (Birkhoff) variety generated by \mathbf{H} . Then the pro- \mathbf{H} topology on $F_{\mathbf{H}}(A)$ is the weakest topology rendering continuous all morphisms of $F_{\mathbf{H}}(A)$ into discrete groups in \mathbf{H} . This topology is induced by the following ultra-norm: for $u \in F_{\mathbf{H}}(A)$, put

$$|u|_{\mathbf{H}} = 2^{-r(u)}, \quad \text{where} \\ r(u) = \min\{|F_{\mathbf{H}}(A)/N| : N \triangleleft F_{\mathbf{H}}(A), u \notin N, F_{\mathbf{H}}(A)/N \in \mathbf{H}\}.$$

Here $r(1)$ is interpreted as ∞ (and hence $|1|_{\mathbf{H}}$ as 0). One then has that $\widehat{F}_{\mathbf{H}}(A)$ is the completion of $F_{\mathbf{H}}(A)$ with respect to the induced metric $d_{\mathbf{H}}(u, v) = |uv^{-1}|_{\mathbf{H}}$. In particular, $F_{\mathbf{H}}(A)$ embeds in $\widehat{F}_{\mathbf{H}}(A)$ in the natural way. Moreover, $\widehat{F}_{\mathbf{H}}(A)$ is the projective limit of the inverse system of all A -generated groups in \mathbf{H} . In other words, $\widehat{F}_{\mathbf{H}}(A) = \varprojlim F_{\mathbf{H}}(A)/N$, where N runs through all normal subgroups of $F_{\mathbf{H}}(A)$ such that $F_{\mathbf{H}}(A)/N \in \mathbf{H}$. Our main focus in this section will be on the Cayley graph $\Gamma_A(\widehat{F}_{\mathbf{H}}(A))$, where A is a finite set. Observe that $\Gamma_A(F_{\mathbf{H}}(A))$ topologically embeds in $\Gamma_A(\widehat{F}_{\mathbf{H}}(A))$ (where $F_{\mathbf{H}}(A)$ is given the pro- \mathbf{H} topology).

Following Gildenhuis and Ribes [17], and Zalesskiĭ and Mel'nikov [59], a profinite graph Γ is *connected* (that is, connected as a profinite graph; the terminology profinitely connected is used in [5], [6], [7], [38]) if each finite continuous morphic image is connected (as an abstract graph). For instance, $\Gamma_A(\widehat{F}_{\mathbf{H}}(A))$ is connected, as it is the projective limit of the Cayley graphs of the finite quotients of $\widehat{F}_{\mathbf{H}}(A)$. By a *subgraph* of a profinite graph we always mean a substructure in the category of profinite graphs, that is, a closed subgraph. We note that, in a profinite graph having only finitely many vertices, a subgraph is connected if and only if it is connected as an abstract graph.

Suppose that Γ is a profinite graph and $v, w \in V(\Gamma)$. Then a connected subgraph Δ of Γ containing v and w is called a *connection* between v and w . The relation on $V(\Gamma)$ of having a connection is clearly a reflexive, symmetric relation. In fact, it is an equivalence relation. This results from the following well-known lemma.

Lemma 2.1. *If Δ, Λ are connected subgraphs of a profinite graph Γ such that $\Delta \cap \Lambda \neq \emptyset$, then $\Delta \cup \Lambda$ is a connected subgraph.*

Proof. The result follows since in any finite image of $\Delta \cup \Lambda$, Δ and Λ map to connected subgraphs with a common vertex. \square

We observe that a profinite graph Γ is connected if and only if any two vertices have a connection. Indeed, if Γ is connected, then Γ is a connection between any

two vertices. Conversely, suppose that $\varphi : \Gamma \rightarrow \Gamma'$ is a morphism with Γ' finite. Let $v, w \in V(\Gamma)$ and Δ be a connection between them. Then $\Delta\varphi$ is a connected subgraph of Γ' containing $v\varphi$ and $w\varphi$. It follows that $\Gamma\varphi$ is connected.

For a profinite ring R with unity, there is a homological notion of what it means for a profinite graph to be an R -tree [17], [59], [6], [5] (see section 9 for more details). For a pseudovariety of groups \mathbf{H} , a graph is said to be an \mathbf{H} -tree if it is an R -tree, where

$$R = \varprojlim \{ \mathbb{Z}/n\mathbb{Z} \mid \mathbb{Z}/n\mathbb{Z} \in \mathbf{H} \}$$

is the free pro-cyclic pro- \mathbf{H} group, considered as a ring. A pseudovariety \mathbf{H} was termed *arborescent* by Almeida and Weil [5], [6] if $\Gamma_A(\widehat{F}_{\mathbf{H}}(A))$ is an \mathbf{H} -tree for all finite sets A . They proved that this holds if and only if \mathbf{H} is non-trivial and closed under co-extension by its Abelian members, that is, $N, G/N \in \mathbf{H}$ and N Abelian implies $G \in \mathbf{H}$. (Throughout, if N is a normal subgroup of G , G is said to be a *co-extension* of G/N by N). Recall that a pseudovariety \mathbf{H} is *extension-closed* if $N, G/N \in \mathbf{H}$ implies $G \in \mathbf{H}$.

A key property of R -trees, $R \neq 0$, is the following (see, for instance, [17], [59], [6]): in an R -tree Γ any pair of vertices v, w has a unique minimal connection (minimality is in terms of containment); this connection is usually denoted by $[v, w]$. This geometric condition is a key ingredient in the papers [5], [6], [52] and in the proof of the Ribes and Zalesskii product theorem [39], [40], [52].

Isolating the aforementioned geometric property of R -trees, we are motivated to define a profinite graph Γ to be *tree-like* if any pair $v, w \in V(\Gamma)$ has a unique minimal connection. We continue to use the notation $[v, w]$ to denote this connection, which we call the *geodesic* between v and w . A tree-like graph must be connected since any two vertices have a connection. We begin with some basic properties of tree-like profinite graphs analogous to the case of R -trees [6], [59]. We note that a finite graph Γ is tree-like if and only if it has no circuits of length greater than one (that is, a finite graph Γ is tree-like if and only if it is a tree, perhaps with loops adjoined to some vertices).

Proposition 2.2. *In any tree-like graph Γ the following hold:*

- (1) *Any connected subgraph of Γ is tree-like.*
- (2) *The intersection $\bigcap_{i \in I} \Gamma_i$ of connected subgraphs Γ_i of Γ is connected.*
- (3) $[u, v] \subseteq [u, w] \cup [w, v]$.
- (4) $[u, v] = [u, w] \cup [w, v]$ if and only if $w \in [u, v]$.
- (5) *Let Γ_1, Γ_2 be connected subgraphs of Γ with $\Gamma_1 \cap \Gamma_2 \neq \emptyset$. Let $v \in V(\Gamma_1)$, $w \in V(\Gamma_2)$. Then $[v, w] \cap \Gamma_1 \cap \Gamma_2 \neq \emptyset$.*
- (6) *If $[u, w] \cap [w, v] = \{w\}$, then $[u, v] = [u, w] \cup [w, v]$.*

Proof. For (1), if $\Delta \subseteq \Gamma$ is a connected, closed subgraph and $v, w \in V(\Delta)$, then $[v, w] \subseteq \Delta$ and is clearly still minimal. For (2), suppose that $\Delta = \bigcap_{i \in I} \Gamma_i \neq \emptyset$ and $v, w \in V(\Delta)$. Then $[v, w] \subseteq \Delta$ since each Γ_i is connected. Thus any two vertices of Δ have a connection. The remaining statements are proved exactly as in the case of R -trees for a profinite ring R [6], [39], [59]. \square

We now aim to prove an important fact about the set of connections between two vertices of a profinite graph. Let $\Gamma = \varprojlim (\Gamma_i, \varphi_{ij})$ be a profinite graph which is the projective limit of an inverse system $(\Gamma_i, \varphi_{ij})_{i, j \in I}$ of finite graphs Γ_i , and denote the canonical projections $\Gamma \rightarrow \Gamma_i$ by π_i . For any profinite graph Δ , let \mathcal{C}_{Δ} be the set of all connected subgraphs of Δ . Each map φ_{ij} induces a map $\widehat{\varphi}_{ij} : \mathcal{C}_{\Gamma_i} \rightarrow \mathcal{C}_{\Gamma_j}$,

$\Lambda \mapsto \Lambda\varphi_{ij}$ for $\Lambda \in \mathcal{C}_{\Gamma_i}$ (which is not, in general, onto). Then $(\mathcal{C}_{\Gamma_i}, \widehat{\varphi}_{ij})$ forms an inverse system and \mathcal{C}_{Γ} can be identified with the profinite set $\varprojlim (\mathcal{C}_{\Gamma_i}, \widehat{\varphi}_{ij})$. So the set of all connected subgraphs of Γ naturally has the structure of a profinite set. Let $\Delta \in \mathcal{C}_{\Gamma}$ and put, for each $i \in I$,

$$\mathcal{U}_i(\Delta) = \{\Lambda \in \mathcal{C}_{\Gamma} \mid \Delta\pi_i = \Lambda\pi_i\}.$$

Then the collection $\{\mathcal{U}_i(\Delta) \mid i \in I\}$ forms a basis for the neighborhoods of Δ in \mathcal{C}_{Γ} (if considered as a profinite set in the aforementioned way). In a similar vein, for $v, w \in V(\Gamma)$, let $\mathcal{C}_{\Gamma}(v, w)$ denote the set of all connections between v and w . Then in the same way, $\mathcal{C}_{\Gamma}(v, w)$ becomes a profinite set, being the projective limit of $\{\mathcal{C}_{\Gamma_i}(v\pi_i, w\pi_i) \mid i \in I\}$; the topology on $\mathcal{C}_{\Gamma}(v, w)$ is described in an analogous way.

Recall that a relation R on a topological space X is said to be *closed* if $R \subseteq X \times X$ is closed in the product topology.

Lemma 2.3. *The containment relation \subseteq on $\mathcal{C}_{\Gamma}(v, w)$ is closed.*

Proof. We show that the complement of \subseteq is open. Let $\mathcal{C} = \mathcal{C}_{\Gamma}(v, w)$ and $\mathcal{C}_i = \mathcal{C}_{\Gamma_i}(v\pi_i, w\pi_i)$ and suppose that $(X, Y) \in \mathcal{C} \times \mathcal{C}$ and X is not contained in Y . Then for some $i \in I$, $X\pi_i$ is not contained in $Y\pi_i$. The induced map $\hat{\pi}_i : \mathcal{C} \times \mathcal{C} \rightarrow \mathcal{C}_i \times \mathcal{C}_i$, $(\Lambda, \Delta) \mapsto (\Lambda\pi_i, \Delta\pi_i)$, is continuous and the latter set is finite. Thus $(X\pi_i, Y\pi_i)\hat{\pi}_i^{-1}$ is a clopen subset of $\mathcal{C} \times \mathcal{C}$ containing (X, Y) , disjoint from \subseteq and hence \subseteq is closed. \square

The following result on closed orders on a compact space is standard.

Lemma 2.4. *Let \leq be a closed partial order on a compact space C . Then, for any element $x \in C$, there exists a minimal element below x .*

Proof. By Zorn's lemma, it suffices to show that any descending chain K of elements below x has a lower bound. For each $k \in K$, let

$$I_k = \{y \in C \mid y \leq k\}.$$

Then I_k is closed since if $(y_j)_{j \in J}$ is a net in C with $y_j \leq k$ for all j and $\lim y_j = y$, then $y \leq k$. Also, the collection $\{I_k\}_{k \in K}$, $k \in K$, clearly satisfies the finite intersection condition, whence $I = \bigcap_{k \in K} I_k \neq \emptyset$. If $y \in I$, then $y \leq k$ for all $k \in K$ as desired. \square

An immediate consequence is the following corollary.

Corollary 2.5. *Let Γ be a profinite graph and $v, w \in V(\Gamma)$. Then any connection Δ between v and w contains a minimal connection.*

We now deduce the following alternate characterization of tree-like graphs.

Proposition 2.6. *Let Γ be a connected profinite graph. Then the following are equivalent:*

- (1) Γ is tree-like;
- (2) given connections Γ_1, Γ_2 between vertices v, w , there is a connection $\Gamma_3 \subseteq \Gamma_1 \cap \Gamma_2$ between v, w ;
- (3) if $\Gamma_1, \dots, \Gamma_n$ are connected subgraphs, then $\Delta = \Gamma_1 \cap \dots \cap \Gamma_n$ is a connected subgraph;
- (4) if $\{\Gamma_i \mid i \in I\}$ is a set of connected subgraphs, then $\Delta = \bigcap_{i \in I} \Gamma_i$ is a connected subgraph.

Proof. Clearly (4) implies (3) implies (2). We have already shown in Proposition 2.2 that (1) implies (4). Thus we are left with showing that (2) implies (1). Let $v, w \in V(\Gamma)$. By Corollary 2.5, with $\Delta = \Gamma$, we see that there is a minimal connection Γ_1 of v and w . Suppose that Γ_2 is any other connection. Then, by (2), $\Gamma_1 \cap \Gamma_2$ contains a connection of v, w which must be Γ_1 by minimality. Thus $[v, w] = \Gamma_1$. The result follows. \square

Next we show that the class of all tree-like graphs is closed under taking projective limits. We prove an auxiliary result first.

Lemma 2.7. *Let Γ and Δ be profinite graphs and $\varphi : \Gamma \rightarrow \Delta$ a morphism. Then, for each pair of vertices $v, w \in \Gamma$, the induced mapping*

$$\hat{\varphi} : \mathcal{C}_\Gamma(v, w) \rightarrow \mathcal{C}_\Delta(v\varphi, w\varphi), \quad \Lambda \mapsto \Lambda\varphi$$

is continuous.

Proof. Let $X \in \mathcal{C}_\Gamma(v, w)$ and let $X' = X\hat{\varphi} = X\varphi \in \mathcal{C}_\Delta(v\varphi, w\varphi)$. Consider a basic neighborhood \mathcal{U} of X' in $\mathcal{C}_\Delta(v\varphi, w\varphi)$. We may assume that there is a finite quotient Δ_i of Δ , with surjective morphism $\psi_i : \Delta \rightarrow \Delta_i$, such that $\mathcal{U} = \{Z \in \mathcal{C}_\Delta(v\varphi, w\varphi) \mid X'\psi_i = Z\psi_i\}$. Then, since $\Gamma\varphi\psi_i$ is a finite quotient of Γ , $\mathcal{V} = \{Y \in \mathcal{C}_\Gamma(v, w) \mid Y\varphi\psi_i = X'\psi_i\}$ is a neighborhood of X in $\mathcal{C}_\Gamma(v, w)$, and by definition $\mathcal{V}\varphi \subseteq \mathcal{U}$. \square

Theorem 2.8. *The class of all tree-like graphs is closed under taking projective limits.*

Proof. Let Γ be the projective limit of the inverse system (Γ_i, φ_{ij}) of tree-like graphs Γ_i . By Proposition 2.2(1), we may assume without loss of generality that the morphisms φ_{ij} are surjective. Denote the canonical projections $\Gamma \rightarrow \Gamma_i$ by φ_i . Choose $v, w \in V(\Gamma)$ and let $\mathcal{C} = \mathcal{C}_\Gamma(v, w)$ and $\mathcal{C}_i = \mathcal{C}_{\Gamma_i}(v\varphi_i, w\varphi_i)$, $i \in I$. For $i \geq j$, let $\widehat{\varphi}_{ij} : \mathcal{C}_i \rightarrow \mathcal{C}_j$ be the induced mapping $\Lambda \mapsto \Lambda\varphi_{ij}$, which is continuous by Lemma 2.7 (but not in general onto). Then $(\mathcal{C}_i, \widehat{\varphi}_{ij})$ forms an inverse system of profinite sets and we may identify \mathcal{C} with the projective limit of that system. (Here we implicitly use that the finite morphic images of Γ are precisely the finite morphic images of the graphs Γ_i .)

We use Proposition 2.6(2) to obtain the result. Let $\Lambda, \Delta \in \mathcal{C}$ and set $\Lambda_i = \Lambda\varphi_i$ and $\Delta_i = \Delta\varphi_i$; then $\Lambda_i, \Delta_i \in \mathcal{C}_i$. For each $i \in I$ let

$$\mathcal{K}_i = \{\Xi_i \in \mathcal{C}_i \mid \Xi_i \subseteq \Lambda_i \cap \Delta_i\}.$$

Each \mathcal{K}_i is a compact subset of \mathcal{C}_i (being the intersection of the two closed order ideals $\{\Xi_i \in \mathcal{C}_i \mid \Xi_i \subseteq \Lambda_i\}$ and $\{\Xi_i \in \mathcal{C}_i \mid \Xi_i \subseteq \Delta_i\}$) and is non-empty since each graph Γ_i is tree-like. Setting $\widehat{\psi}_{ij} = \widehat{\varphi}_{ij}|_{\mathcal{K}_i}$, we obtain the inverse system $(\mathcal{K}_i, \widehat{\psi}_{ij})$. Now we have

$$\emptyset \neq \varprojlim \mathcal{K}_i \subseteq \varprojlim \mathcal{C}_i = \mathcal{C}.$$

By construction, each member Ξ of $\varprojlim \mathcal{K}_i$, viewed as an element of \mathcal{C} , is a connection in Γ between v and w and is contained in $\Lambda \cap \Delta$ (since, for all i , $\Xi\varphi_i$ is contained in $\Lambda_i \cap \Delta_i$ and is a connection between $v\varphi_i$ and $w\varphi_i$). It follows that Γ is tree-like. \square

We define a pseudovariety of groups \mathbf{H} to be *arboreous* if the Cayley graphs $\Gamma_A(\widehat{F}_{\mathbf{H}}(A))$ are tree-like for all finite sets A . Any arborescent pseudovariety is arboreous. Clearly the trivial pseudovariety is, as well (since if \mathbf{H} contains only the trivial group, then $\widehat{F}_{\mathbf{H}}(A)$ is the trivial group and hence its Cayley graph has

a single vertex). We shall see that the class of arboreous pseudovarieties is much bigger than this, and enjoys many nice properties. Two immediate consequences of Theorem 2.8 are the following.

Corollary 2.9. *For each arboreous pseudovariety \mathbf{H} and each profinite set A , $\Gamma_A(\widehat{F_{\mathbf{H}}}(A))$ is tree-like.*

Proof. If $A = \varprojlim A_i$, then $\widehat{F_{\mathbf{H}}}(A) = \varprojlim \widehat{F_{\mathbf{H}}}(A_i)$ and hence

$$\Gamma_A(\widehat{F_{\mathbf{H}}}(A)) = \varprojlim \Gamma_{A_i}(\widehat{F_{\mathbf{H}}}(A_i)). \quad \square$$

For the next result we take into account that if \mathbf{H} is the union of the upwardly directed set $\{\mathbf{H}_i \mid i \in I\}$, then, for each finite set A , $\widehat{F_{\mathbf{H}}}(A) = \varprojlim \widehat{F_{\mathbf{H}_i}}(A)$.

Corollary 2.10. *The set of all arboreous pseudovarieties of groups is closed under taking unions of upwardly directed sets.*

Next we state some easy necessary conditions for a pseudovariety to be arboreous. If A is a set, we use $FG(A)$ to denote the free group on A .

Proposition 2.11. *Let \mathbf{H} be a non-trivial arboreous pseudovariety. Then, for any finite set A , $FG(A)$ is residually \mathbf{H} .*

Proof. Let \mathbf{H} be non-trivial. We have already observed that $\Gamma_A(F_{\mathbf{H}}(A))$ embeds in $\Gamma_A(\widehat{F_{\mathbf{H}}}(A))$. If $FG(A)$ were not residually \mathbf{H} (that is, $FG(A) \neq F_{\mathbf{H}}(A)$), then $\Gamma_A(F_{\mathbf{H}}(A))$ would have a minimal circuit of length greater than one. This circuit would be a connected, closed subgraph and hence tree-like by Proposition 2.2(1). But a finite graph with a circuit of length at least two cannot not be tree-like — a contradiction. \square

The above proof motivates the following straightforward fact; see [48] for the arborescent case.

Proposition 2.12. *Let $g, h \in FG(A) \subseteq \widehat{F_{\mathbf{H}}}(A)$ with \mathbf{H} non-trivial, arboreous and A finite. Then $[g, h]$ is the usual geodesic between g and h in $\Gamma_A(FG(A))$.*

Proof. Clearly the usual geodesic X of $\Gamma_A(FG(A))$ gives a connection between g and h ; any smaller connection would then exist within $\Gamma_A(FG(A))$. \square

Recall that if \mathbf{H}_1 and \mathbf{H}_2 are pseudovarieties, then their join $\mathbf{H} = \mathbf{H}_1 \vee \mathbf{H}_2$ is the smallest pseudovariety containing them both. It is well known that, for any finite set A , the diagonal embedding of A provides a representation of $\widehat{F_{\mathbf{H}}}(A)$ as a subdirect product of $\widehat{F_{\mathbf{H}_1}}(A)$ and $\widehat{F_{\mathbf{H}_2}}(A)$. A pseudovariety \mathbf{H} is *join irreducible* if $\mathbf{H} = \mathbf{H}_1 \vee \mathbf{H}_2$ implies that $\mathbf{H}_i = \mathbf{H}$ for some i . The following result uses Theorem 3.7, proved below.

Corollary 2.13. *Each arboreous pseudovariety is join irreducible.*

Proof. Suppose that \mathbf{H} is arboreous and $\mathbf{H} = \mathbf{H}_1 \vee \mathbf{H}_2$ with $\mathbf{H}_1, \mathbf{H}_2 \neq \mathbf{H}$. Then there exist finite sets A_1, A_2 (which we may take to be disjoint) and elements $\pi_1 \in \widehat{F_{\mathbf{H}_1}}(A_1) \setminus \{1\}$, $\pi_2 \in \widehat{F_{\mathbf{H}_2}}(A_2) \setminus \{1\}$ such that π_1 maps to 1 in $\widehat{F_{\mathbf{H}_1}}(A_1)$ and π_2 maps to 1 in $\widehat{F_{\mathbf{H}_2}}(A_2)$. Let $A = A_1 \cup A_2$. Using the subdirect product embedding referred to above, we see that $\pi_1 \pi_2 = \pi_2 \pi_1$ in $\widehat{F_{\mathbf{H}}}(A)$, or equivalently $\pi_1 \pi_2 \pi_1^{-1} \pi_2^{-1} = 1$. But, from Theorem 3.7, it follows immediately that $\pi_1 \pi_2 \pi_1^{-1} \pi_2^{-1} \neq 1$ because $[1, \pi_1]$ and $[1, \pi_2]$ have disjoint sets of labels (the former being labeled over A_1 , the latter over A_2), whence the hypotheses of that theorem are clearly satisfied. \square

Thus \mathbf{G}_{nil} is not arboreous. Since extension-closed pseudovarieties are arborescent and hence arboreous, it follows that \mathbf{G} , \mathbf{G}_p , and \mathbf{G}_{sol} are all join irreducible.

3. DENDRAL GROUPS

This section can be read independently of the remainder of the paper (outside of notation and definitions). We define a profinite group G to be *dendral with respect to a profinite generating set* A if the Cayley graph $\Gamma_A(G)$ is tree-like. More generally, a profinite group will be termed *dendral* if it is dendral with respect to some profinite generating set. In this section, we present two purely group-theoretic results. The first is a Bass-Serre type characterization of dendral groups: they are precisely those groups admitting a free action on a tree-like graph with finite vertex quotient. As a consequence, it follows that each open subgroup of a dendral group is again dendral. The second result is the cancellation result mentioned in the Introduction. This latter result will be of essential use in section 5. In this section, abstract groups are allowed to be infinite.

3.1. A Bass-Serre type theorem. Let G be a profinite group, A a profinite set and $\psi : A \rightarrow G$ an arbitrary continuous map ($A\psi$ not necessarily generating G). Then the definition of $\Gamma_A(G)$ can be extended in an obvious way. The following is [17, Proposition 1.1], and is easily deduced from the discrete case [44].

Proposition 3.1. *The profinite set A generates G (via ψ) if and only if $\Gamma_A(G)$ is connected.*

Throughout, all graphs to be considered are oriented. Following Serre, we say that a group acts freely on a graph Γ if it acts without fixed points on the vertex set and preserves the orientation of the edges. For instance, the left action of a group on its Cayley graph is free.

We first recall the proof, given in [44], that an abstract group is free if and only if it acts freely on a tree. We shall modify the definitions and arguments slightly in order to suit our purposes; in particular, so that they can be transferred to the profinite setting.

Let Γ be an abstract graph and $\Lambda = \{T_i \mid i \in I\}$ a collection of pairwise disjoint subtrees of Γ . Without loss of generality, we assume that each vertex of Γ is contained in some T_i . If, for some family Λ of subtrees, this were not the case, we could always add one-vertex subtrees to the family to meet the assumption. We define the graph Γ/Λ by the following data: $V(\Gamma/\Lambda) = \Lambda$, that is, the vertices of Γ/Λ are just the elements T_i of the family Λ ; the edges of Γ/Λ are just the edges of Γ ; the incidence mappings $\iota', \tau' : E(\Gamma/\Lambda) \rightarrow V(\Gamma/\Lambda)$ are defined by $\iota'(e) = T_i$ if and only if $\iota(e) \in T_i$, and analogously for τ' . The mapping $e \mapsto e^{-1}$ remains unchanged. Note that the edges of T_i in Γ/Λ become loops at the vertex T_i .

There is a natural graph morphism $\varphi : \Gamma \rightarrow \Gamma/\Lambda$ where, for $v \in V(\Gamma)$, $v\varphi = T_i \in V(\Gamma/\Lambda)$ if and only if $v \in T_i$ and φ is the identity mapping on the set of edges. Clearly, Γ/Λ is connected if and only if Γ is connected. Moreover, the inverse image in Γ of any connected subgraph of Γ/Λ that contains all loops of Λ at each of its vertices is connected. If Γ is a tree, then Γ/Λ is a tree with loops adjoined at those vertices v of Γ/Λ which correspond to the non-trivial subtrees T_i of Λ [44]. We say that Γ/Λ is obtained from Γ by *contracting the trees* of Λ . The difference between our definition of Γ/Λ and that in [44] is that, in the latter case, all edges which

come from edges in non-trivial trees T_i (which become loops in our definition) are removed from the graph.

Suppose that G acts freely on Γ and T' is a maximal tree in $G \backslash \Gamma$ (which denotes the quotient graph formed from the G -orbits). Then it is well known [44] that there is a subtree T in Γ mapping isomorphically to T' under the quotient mapping $\Gamma \rightarrow G \backslash \Gamma$. Let Γ' be the connected subgraph of Γ consisting of T and all positive edges with initial vertex in T (and inverses of such edges). Observe that no two edges of Γ' can be identified under the quotient mapping. Indeed, if two edges e_1, e_2 are identified, we may assume that they are both positive. Also $\iota e_1, \iota e_2 \in T$, so we must have $\iota e_1 = \iota e_2$. So if $g e_1 = e_2$, then $g \iota e_1 = \iota e_2 = \iota e_1$, whence $g = 1$. Thus $e_1 = e_2$. Moreover, each edge of $G \backslash \Gamma$ has a (unique) lift to Γ' since any positive edge of $G \backslash \Gamma$ has initial vertex in T' and hence a lift with initial vertex in T .

The tree T is called a *tree of representatives* and the graph Γ' is the associated *fundamental G -transversal*. Note that $gT \cap g'T = \emptyset$ whenever $g \neq g'$, so $G \cdot T$ is a disjoint union of subtrees of Γ and $G \cdot T$ contains all vertices of Γ , that is, $V(\Gamma) = V(G \cdot T)$. Consequently, for each positive edge e in Γ' , there is a unique element $t_e \in G$ such that $\tau e \in t_e T$, whence $t_e^{-1} \tau e \in T$.

Let A be the set of positive edges of Γ' . This will be our alphabet: the *edges* of the graph Γ' will be the *labels* of the edges of the Cayley graph of G . The mapping $t : A \rightarrow G$ given by $e \mapsto t_e$ allows us to form the Cayley graph $\Gamma_A(G)$. The essential step now is to show that the graphs $\Gamma/(G \cdot T)$ and $\Gamma_A(G)$ are isomorphic. In the abstract case, this will then suffice to show that if Γ is a tree, then $\Gamma_A(G)$ is a tree with loops adjoined corresponding to those generators e with $e \in T$; such e represent 1 in G since, for $e \in T$, $t_e = 1$. One could remove all such generators, that is, consider instead of A the set of positive edges of Γ' which are not in T . For the graph $\Gamma/(G \cdot T)$, this would correspond to removing all edges of T when we contract the trees gT to a point, instead of converting them to loops (as we are doing here). But, for our purposes, we need to retain these edges.

The isomorphism between $\Gamma/(G \cdot T)$ and $\Gamma_A(G)$ can be realized as follows. We use (gT) to denote the vertex of $\Gamma/(G \cdot T)$ corresponding to gT . As the trees gT ($g \in G$) are pairwise disjoint, the map $g \mapsto (gT)$ is a bijection, the inverse map gives a bijection from $V(\Gamma/(G \cdot T))$ to $V(\Gamma_A(G))$. We extend the map to edges. It suffices to deal with positive edges. Let e be a positive edge from (gT) to (hT) . Then e is a positive edge of Γ linking gT with hT , whence $g^{-1}e \in \Gamma'$ (and is also a positive edge) and $\tau(g^{-1}e) \in g^{-1}hT$. We send e to $(g, g^{-1}e)$. To see that this is a graph morphism, it suffices to show that $gt_{g^{-1}e} = h$. But this is clear from the definition since $\tau(g^{-1}e) \in g^{-1}hT$ and hence $t_e = g^{-1}h$. Let e, f be distinct edges of $\Gamma/(G \cdot T)$ (that is, edges of Γ). Then either they belong to distinct G -orbits, in which case they map to edges having distinct labels, or $e = gf$ with $g \neq 1$ — in this case they map to edges having distinct initial vertices. Hence the mapping is injective on the set of edges. Finally, for e a positive edge in Γ' and $g \in G$, ge is a positive edge from gT to $gt_e T$, whence ge maps to (g, e) , so the map is also surjective. Therefore, the following result is proved (the arguments are essentially those of [44]).

Theorem 3.2. *Let G be a group acting freely on a graph Γ with a tree of representatives T and associated fundamental G -transversal Γ' . Let A be the set of positive edges of Γ' and $t : A \rightarrow G$ be defined by $\tau e \in t_e T$. Then $\Gamma/(G \cdot T) \cong \Gamma_A(G)$, the*

Cayley graph being defined via the mapping t . In particular, if Γ is connected, then A generates G . Moreover, if Γ is a tree, then $\Gamma_A(G)$ is a tree (and so G is free).

Our goal is to develop a profinite analog of this result. In the context of profinite groups and profinite graphs, all actions shall be assumed continuous. We begin by stating an important result of Zalesskiĭ [58, Lemma 1.3]; see also [43].

Proposition 3.3. *Suppose that a profinite group G acts freely on a profinite graph Γ . Then there are inverse systems $\{G_i \mid i \in I\}$ of finite groups and $\{\Gamma_i \mid i \in I\}$ of finite graphs for a common index set I such that $G = \varprojlim G_i$ and $\Gamma = \varprojlim \Gamma_i$; moreover, for each $i \in I$, G_i acts freely on Γ_i , these actions commute with the morphisms of the inverse systems, and the action of G on Γ is induced by these actions.*

Suppose that G has a free action on Γ such that $G \backslash \Gamma$ has a finite vertex set; let G_i and Γ_i be as per Proposition 3.3. Without loss of generality, we may assume that the inverse systems consist of surjective maps. A straightforward compactness argument shows that $G \backslash \Gamma = \varprojlim G_i \backslash \Gamma_i$. Also, the trees T' and T (as above) are finite. By going to a cofinal subset of I , we may therefore assume that $T \cong T'$ embeds in $G_i \backslash \Gamma_i$ for all i . That is, each Γ_i contains a subtree T_i isomorphic with T and the trees T_i are isomorphically mapped onto each other under the morphisms of the inverse system. Moreover, each quotient graph $G_i \backslash \Gamma_i$ contains a subtree T'_i (each isomorphic with T), and T_i is isomorphically mapped to T'_i under the quotient mapping $\Gamma_i \rightarrow G_i \backslash \Gamma_i$.

As above, let Γ'_i be the subgraph of Γ_i consisting of T_i and of the positive edges going out of T_i (and their inverses), and let Γ' be the analogous subgraph of Γ . Then we have $\Gamma' = \varprojlim \Gamma'_i$. Let A , respectively A_i , be the set of positive edges of Γ' , respectively Γ'_i . Then A is a profinite set and $A = \varprojlim A_i$. The profinite set A shall be our profinite generating set. Here one can see the reason for the modifications we mentioned earlier: an edge in T may be a cluster point in the set of edges of Γ' , so we cannot guarantee that we still have a profinite set if we remove the edges of T from A .

Now, for each i , we have a natural mapping $t^{(i)} : A_i \rightarrow G_i$ defined in the same way as above (that is, for $e \in A_i$, $\tau e \in t_e^{(i)} T_i$), this mapping gives rise to the Cayley graph $\Gamma_{A_i}(G_i)$. Likewise we have the map $t : A \rightarrow G$, which is continuous because $A = \varprojlim A_i$, the mappings $t^{(i)} : A_i \rightarrow G_i$ commute with the morphisms of the inverse system, and t is easily verified to be the mapping induced by the mappings $t^{(i)}$, $i \in I$. Moreover, t gives rise to the profinite Cayley graph $\Gamma_A(G)$. From Theorem 3.2, we have $\Gamma_i/(G_i \cdot T_i) \cong \Gamma_{A_i}(G_i)$ for all i , and $\Gamma_A(G) = \varprojlim \Gamma_{A_i}(G_i)$. Since all the graphs $\Gamma_i/(G_i \cdot T_i)$ are connected, all the $\Gamma_{A_i}(G_i)$ are connected and thus $\Gamma_A(G)$ is connected as a profinite graph. Therefore A generates G . From Proposition 3.3, it also follows that $\Gamma/(G \cdot T) = \varprojlim \Gamma_i/(G_i \cdot T_i)$, whence $\Gamma/(G \cdot T) \cong \Gamma_A(G)$. The natural projection $\Gamma \rightarrow \Gamma/(G \cdot T)$ is continuous because all the projections $\Gamma_i \rightarrow \Gamma_i/(G_i \cdot T_i)$ commute with the morphisms of the inverse systems, and the projection $\Gamma \rightarrow \Gamma/(G \cdot T)$ is induced by these latter projections. In particular, $\Gamma_A(G)$ is a morphic image of Γ , obtained by contracting the trees $\{gT \mid g \in G\}$.

Let us, for the moment, call a subgraph Δ of a graph Λ *loop complete* if Δ contains all loops of Λ at the vertices of Δ . The projections $\Gamma_i \rightarrow \Gamma_i/(G_i \cdot T_i)$ have the property that the inverse image of each loop complete, connected subgraph is connected (this was mentioned earlier when we discussed the abstract case).

Consequently, for the projection $\Gamma \rightarrow \Gamma/(G \cdot T)$ we also have that the inverse image of each loop complete, connected subgraph of $\Gamma/(G \cdot T)$ is connected in Γ .

It remains to show that if Γ is tree-like, then $\Gamma_A(G)$ is as well. Indeed, let g, h be vertices of $\Gamma_A(G)$ and let Ξ be a connection between g and h . Let $[\Xi]$ be the (abstract) graph obtained from Ξ by adjoining all loops. Then $[\Xi]$ is a loop complete, connected profinite subgraph of $\Gamma_A(G)$. Indeed, $V([\Xi]) = V(\Xi)$ is closed. Suppose that $\{e_\alpha\}$ is a net of elements of $E([\Xi])$ converging to e . If e is a loop, then $e \in E([\Xi])$ by construction. If e is not a loop, then, by going to a cofinal subnet, we may assume that none of the edges e_α is a loop and thus $\{e_\alpha\} \subseteq E(\Xi)$. Hence $e \in E(\Xi) \subseteq E([\Xi])$ and $E([\Xi])$ is closed. Since Ξ is connected and $V([\Xi]) = V(\Xi)$, it is clear that $[\Xi]$ is connected.

Now consider the inverse image $[\Xi]\psi^{-1}$ in Γ where $\psi : \Gamma \rightarrow \Gamma/(G \cdot T) \cong \Gamma_A(G)$ is the natural projection. Clearly $gT \cup hT \subseteq [\Xi]\psi^{-1}$. Let Δ be the least connected (profinite) subgraph of Γ containing the set $gT \cup hT$ (Γ is tree-like, so Δ exists and is well defined by Proposition 2.2(2)). Then $\Delta \subseteq [\Xi]\psi^{-1}$. Consequently, $\Delta\psi \subseteq [\Xi]$. Let $(\Delta\psi)$ be the (abstract) graph obtained from $\Delta\psi$ by removing all loops. Then $(\Delta\psi) \subseteq \Xi$ and so $\overline{(\Delta\psi)} \subseteq \Xi$. But $\overline{(\Delta\psi)}$ is a connected profinite graph containing g and h . Hence each connection between g and h contains the connection $\overline{(\Delta\psi)}$. Therefore $\Gamma_A(G)$ is tree-like.

We have thus proved a profinite analog of Theorem 3.2, namely:

Theorem 3.4. *Let G be a profinite group acting freely on a profinite graph Γ with finite vertex quotient. Let T be a tree of representatives, Γ' be the associated fundamental G -transversal, A be the set of positive edges of Γ' , and $t : A \rightarrow G$ be defined by $\tau e \in t_e T$. Then t is continuous, and $\Gamma/(G \cdot T)$ and $\Gamma_A(G)$ are isomorphic profinite graphs (the Cayley graph being defined via the mapping t). In particular, if Γ is connected, then A generates G . Moreover, if Γ is tree-like, then $\Gamma_A(G)$ is a tree-like (and so G is dendral).*

Observing that the quotient by the natural (free) left action of a group on its Cayley graph has a single vertex, we see that we have proved the following.

Theorem 3.5. *A profinite group G is dendral if and only if it admits a free (continuous) action on a tree-like graph Γ with finite vertex quotient.*

Of particular interest is the following special case.

Corollary 3.6. *Each open subgroup of a dendral group is dendral.*

Proof. An open subgroup H of a profinite group G is profinite and of finite index. Suppose that $\Gamma_A(G)$ is tree-like. Then H has a free action on $\Gamma_A(G)$ with finite vertex quotient. \square

The authors are grateful to P. Zalesskiĭ for pointing out Proposition 3.3, which improved the original version of the result considerably.

3.2. Cancellation in dendral groups. Here we present a cancellation result for dendral groups, a result which is a natural analog of a well-known and obvious fact holding in free groups. This result will be of essential use in section 5.

Suppose that g_1, g_2, \dots, g_n belong to some free group $FG(A)$ where $n \geq 3$. Viewing the elements g_i as reduced words over \tilde{A}^* , a sufficient condition to guarantee that $g_1 \cdots g_n \neq 1$ is, for all $i = 2, \dots, n-1$, a letter of g_i , remains uncanceled in

the reduced form of $g_{i-1}g_i g_{i+1}$. Let us translate this condition geometrically. As the Cayley graph $\Gamma_A(FG(A))$ is a tree, we continue to use $[f, g]$, $f, g \in FG(A)$, to denote the unique geodesic in $\Gamma_A(FG(A))$ between f and g . The condition that a letter of g_i remains uncanceled in the reduced form of $g_{i-1}g_i g_{i+1}$ is easily seen to be equivalent to

$$[1, g_{i-1}] \cap [g_{i-1}, g_{i-1}g_i] \cap [g_{i-1}g_i, g_{i-1}g_i g_{i+1}] = \emptyset.$$

While “reduced form” has no obvious meaning in profinite groups, the above equality does have an interpretation in free profinite groups and, more generally, in dendral groups. For such groups, in particular, for finitely generated free pro- \mathbf{H} -groups of an arboreous pseudovariety \mathbf{H} , a precise analog of the above-mentioned fact holds. It may be formulated as follows:

Theorem 3.7. *Let G be an A -generated dendral group with A a profinite set. Let $\pi_1, \dots, \pi_n \in G$, $n \geq 3$. If, for each $k = 2, \dots, n - 1$,*

$$(3.1) \quad [1, \pi_{k-1}] \cap [\pi_{k-1}, \pi_{k-1}\pi_k] \cap [\pi_{k-1}\pi_k, \pi_{k-1}\pi_k\pi_{k+1}] = \emptyset,$$

then $\pi_1 \cdots \pi_n \neq 1$.

We decompose the proof into several lemmas. Throughout, G is assumed to be a profinite group, which is dendral with respect to a profinite generating set A . Let $\pi, \sigma \in G$; we know that $[1, \pi\sigma] \subseteq [1, \pi] \cup \pi[1, \sigma]$. The edges in $([1, \pi] \cup \pi[1, \sigma]) \setminus [1, \pi\sigma]$ are said to be *cancelled* when multiplying π by σ ; the edges in $[1, \pi\sigma]$ are said to *remain uncanceled*. The first two lemmas are easy, but useful observations which give sufficient conditions to ensure that edges remain uncanceled. The first states that cancelled edges must be in the intersection of appropriate graphs. We state the results in their most general form. Throughout, vertices are denoted by Greek letters α, β , etc., while edges are represented as pairs $(\beta, b), (\gamma, c)$, etc.

Lemma 3.8. (1) *If $(\gamma, c) \in [\alpha, \pi]$ but $(\gamma, c) \notin [\alpha, \pi\sigma]$, then $(\gamma, c) \in [\alpha, \pi] \cap \pi[1, \sigma]$.*
 (2) *If $(\beta, b) \in \pi[1, \sigma]$ but $(\beta, b) \notin [\gamma, \pi\sigma]$, then $(\beta, b) \in [\gamma, \pi] \cap \pi[1, \sigma]$.*

Proof. (1) By assumption, $(\gamma, c) \in [\alpha, \pi] \subseteq [\alpha, \pi\sigma] \cup [\pi\sigma, \pi] = [\alpha, \pi\sigma] \cup \pi[1, \sigma]$; if (γ, c) is not in $[\alpha, \pi\sigma]$, then $(\gamma, c) \in \pi[1, \sigma]$. Statement (2) is proved dually. \square

The next statement gives a sufficient condition in order for an edge to be protected from cancellation; when reading the proof it may be useful to draw a picture.

Lemma 3.9. *Let $(\beta_1, b_1) \in [\alpha, \pi]$; if there exists an edge $(\gamma_1, c_1) \in [\beta_1, \pi] \setminus \pi[1, \sigma]$, then $(\beta_1, b_1) \in [\alpha, \pi\sigma]$.*

Proof. We have $(\beta_1, b_1) \in [\alpha, \pi] \subseteq [\alpha, \pi\sigma] \cup [\pi\sigma, \pi]$; if $(\beta_1, b_1) \notin [\alpha, \pi\sigma]$, then $(\beta_1, b_1) \in [\pi, \pi\sigma] = \pi[1, \sigma]$. But then $(\gamma_1, c_1) \in [\beta_1, \pi] \subseteq [\pi, \pi\sigma] = \pi[1, \sigma]$, a contradiction to our assumption. \square

Notice that, in the context of Lemma 3.9, we also have $(\gamma_1, c_1) \in [\beta_1, \pi\sigma]$ by Lemma 3.8(1) with $\alpha = \beta_1$. The dual of Lemma 3.9 is also true, but will not be used. We intend to show that, for $\pi_1, \pi_2, \dots, \pi_n \in G$ satisfying (3.1), the graphs $[1, \pi_1\pi_2], \dots, [1, \pi_1 \cdots \pi_n]$ never have empty edge sets. The idea of the proof is to use the above two lemmas in alternation. Lemma 3.9 guarantees that, under its hypothesis, the membership of (γ, c_1) in $[\beta_1, \pi_1]$ protects any edge (β_1, b_1) from

being cancelled against π_2 ; Lemma 3.8(2) is aimed at establishing a situation as in the hypothesis of Lemma 3.9 for the next step.

We continue to assume that G is a profinite group generated by a profinite set A such that $\Gamma_A(G)$ is tree-like. For convenience we identify A with its image in G (we note that distinct generators, which are not sent to 1, are necessarily sent to distinct elements of G). Suppose that $G = \varprojlim G_i$ for some inverse quotient system $(G_i, \varphi_{ij})_{i,j \in I}$ of finite groups. We denote the canonical projections $G \rightarrow G_i$ by φ_i . Note that $A\varphi_i$ is a generating set of G_i , so we may consider the Cayley graph $\Gamma_A(G_i)$ of G_i with respect to A — this is a profinite graph having only finitely many vertices. So, for each pair v, w of vertices of $\Gamma_A(G_i)$, the set $\iota^{-1}v \cap \tau^{-1}w$ of all edges going from v to w is clopen. Abusing notation, we use φ_i , respectively φ_{ij} , for the graph morphisms $\Gamma_A(G) \rightarrow \Gamma_A(G_i)$, respectively $\Gamma_A(G_i) \rightarrow \Gamma_A(G_j)$ (for $i \geq j$) induced by φ_i , respectively φ_{ij} , as well as for the induced graph morphisms between subgraphs of these graphs.

Let $\pi, \rho, \tau \in G$ and set $P = [1, \pi]$, $R = \pi[1, \rho]$ and $T = \pi\rho[1, \tau]$; put $M = P \cap R$ and $N = R \cap T$. For each $i \in I$ and each $U \in \{P, R, T, M, N\}$, put $U_i = U\varphi_i$; then, for any $i \geq j$, $U_j = U_i\varphi_{ij}$ and $(U_i, \varphi_{ij})_{i,j \in I}$ forms an inverse system of graphs with $U = \varprojlim U_i$. Next, for each $i \in I$, let X_i be the largest connected subgraph of $P_i \cap R_i$ containing $\pi\varphi_i$. (Note: a vertex v belongs to this graph if and only if there is a path in $P_i \cap R_i$ (in the usual sense) from $\pi\varphi_i$ to v and an edge is in X_i if it is in $P_i \cap R_i$ and either its initial or its terminal vertex can be reached from $\pi\varphi_i$ by a path in $P_i \cap R_i$ — one easily verifies that X_i , so constructed, is closed; hence the graph X_i is well defined.) Likewise, define Y_i to be the largest connected subgraph of $R_i \cap T_i$ containing $\pi\rho\varphi_i$. Then $M_i \subseteq X_i$ and $N_i \subseteq Y_i$ for all i . Moreover, for all $i \geq j$, $X_i\varphi_{ij} \subseteq X_j$ and $Y_i\varphi_{ij} \subseteq Y_j$. So (X_i, φ_{ij}) and (Y_i, φ_{ij}) form inverse systems of graphs (although the morphisms φ_{ij} need not be surjective here). From $M_i \subseteq X_i$ and $N_i \subseteq Y_i$ for all i , we get $M = \varprojlim M_i \subseteq \varprojlim X_i$ and $N = \varprojlim N_i \subseteq \varprojlim Y_i$. But $\varprojlim X_i$ is a connected subgraph of $P \cap R = M$ (since $X_i \subseteq P_i$ and $X_i \subseteq R_i$ for all i), so $M = \varprojlim X_i$. Likewise $N = \varprojlim Y_i$. If $X_i \cap Y_i \neq \emptyset$, then $X_j \cap Y_j \supseteq (X_i \cap Y_i)\varphi_{ij} \neq \emptyset$ for all $j \leq i$; in other words, if $X_j \cap Y_j = \emptyset$, then $X_i \cap Y_i = \emptyset$ for all $i \geq j$. Suppose that $X_i \cap Y_i \neq \emptyset$ for all i . Set $Z_i = X_i \cap Y_i$; then $Z_i\varphi_{ij} \subseteq Z_j$ for all $i \geq j$, and again (Z_i, φ_{ij}) forms an inverse system of non-empty graphs, whence $\emptyset \neq \varprojlim Z_i \subseteq \varprojlim X_i \cap \varprojlim Y_i = M \cap N$. We have obtained the following lemma.

Lemma 3.10. *If $M \cap N = \emptyset$, then there exists $i \in I$ such that $X_j \cap Y_j = \emptyset$ for all $j \geq i$.*

Now let $n \geq 3$ and $\pi_1, \dots, \pi_n \in G$; for $k = 1, \dots, n$, set

$$A_k = \pi_1 \dots \pi_{k-1}[1, \pi_k]$$

and, for $k = 1, \dots, n-1$, let

$$M_k = A_k \cap A_{k+1}.$$

For each $k = 1, \dots, n$ and each $i \in I$, let $A_{ki} = A_k\varphi_i$. Then A_{ki} is a connected subgraph of $\Gamma_A(G_i)$ containing $(\pi_1 \dots \pi_{k-1})\varphi_i$ and $(\pi_1 \dots \pi_k)\varphi_i$. For each $k = 1, \dots, n-1$, let X_{ki} be the largest connected subgraph of $A_{ki} \cap A_{k+1,i}$ containing $(\pi_1 \dots \pi_k)\varphi_i$. The above arguments show that (X_{ki}, φ_{ij}) is an inverse system of graphs and $\varprojlim_i X_{ki} = M_k$. From Lemma 3.10 we obtain the next result.

Corollary 3.11. *Suppose that $M_k \cap M_{k+1} = \emptyset$ for all $k = 1, \dots, n - 2$; then there exists $i \in I$ such that $X_{kj} \cap X_{k+1j} = \emptyset$ for all $j \geq i$ and all $k = 1, \dots, n - 2$.*

We continue to use the notation introduced above. Assume that $M_k \cap M_{k+1} = \emptyset$ for all $k = 1, \dots, n - 2$. By Corollary 3.11 and by going to a cofinal subset, we may assume that $X_{ki} \cap X_{k+1i} = \emptyset$ for all $k = 1, \dots, n - 2$ and all $i \in I$.

The following definitions are crucial. For $k = 1, \dots, n - 1$ and $i \in I$, let

$$C_{ki} = \{(g, c) \in A_{ki} \mid g \notin X_{ki}, gc \in X_{ki}\}$$

and, for $k = 2, \dots, n - 1$ and $i \in I$, let

$$B_{ki} = \{(h, b) \in A_{ki} \mid h \in X_{k-1i}, hb \notin X_{k-1i}\}.$$

For all $i \in I$ and $k = 2, \dots, n - 1$, C_{ki} and B_{ki} are non-empty (this follows from $A_{ki} \setminus X_{ki} \neq \emptyset$, respectively, $A_{ki} \setminus X_{k-1i} \neq \emptyset$, which in turn follows from $X_{k-1i} \cap X_{ki} = \emptyset$). Also $C_{ki} \cap A_{k+1i} = \emptyset = B_{ki} \cap A_{k-1i}$ (this follows from the definition of X_{ki} , respectively X_{k-1i} , as maximal connected subgraphs). In addition, $B_{ki} \cap X_{ki} = \emptyset$ (this again follows from $X_{k-1i} \cap X_{ki} = \emptyset$); similarly, $C_{k+1i} \cap X_{ki} = \emptyset$. (These latter two equations hold for those k for which they make sense.)

Next, for $k = 1, \dots, n - 1$, set

$$C_k = \{(\gamma, c) \in A_k \mid (\gamma, c)\varphi_i \in C_{ki} \text{ for some } i \in I\}$$

and for $k = 2, \dots, n - 1$ let

$$B_k = \{(\eta, b) \in A_k \mid (\eta, b)\varphi_i \in B_{ki} \text{ for some } i \in I\}.$$

Then $C_k \cap A_{k+1} = B_k \cap A_{k-1} = B_k \cap M_k = C_{k+1} \cap M_k = \emptyset$ for all k such that these expressions make sense. Indeed, suppose that $(\eta, e) \in C_k \cap A_{k+1}$; then $(\eta, e)\varphi_i \in C_{ki}$ for some $i \in I$ and $(\eta, e)\varphi_i \in A_{k+1i}$, but $C_{ki} \cap A_{k+1i} = \emptyset$, a contradiction; similar arguments apply to the other cases. All the sets C_k and B_k are non-empty ($k = 2, \dots, n - 1$) because the corresponding sets C_{ki} and B_{ki} are non-empty.

The intuitive idea behind these definitions is the following. For each $i \in I$ and k , the set C_{ki} forms a “barrier” of edges that must be passed if one travels inside A_{ki} from $(\pi_1 \cdots \pi_{k-1})\varphi_i$ to $(\pi_1 \cdots \pi_k)\varphi_i$ (and similarly for B_{ki}) and, moreover, no edge of C_{ki} occurs in A_{k+1i} . In the limit, this guarantees that the edges of C_k protect the edges in $A_k \setminus M_k$ from being cancelled when π_k is multiplied by π_{k+1} on the right.

Lemma 3.12. *For each vertex $\varepsilon \in A_k \setminus M_k$, there is an edge $(\gamma, c) \in C_k$ such that $(\gamma, c) \in [\varepsilon, \pi_1 \cdots \pi_k]$ (here $1 \leq k \leq n - 1$).*

Proof. Note that the arguments before Lemma 3.10 show that $M_k = \varprojlim_i X_{ki}$. Hence there exists $i \in I$ such that $\varepsilon\varphi_i \notin X_{ki}$. Observe that $[\varepsilon, \pi_1 \cdots \pi_k]\varphi_i$ is a connected subgraph of A_{ki} containing $\varepsilon\varphi_i$ and $(\pi_1 \cdots \pi_k)\varphi_i$ such that $\varepsilon\varphi_i \notin X_{ki}$, but $(\pi_1 \cdots \pi_k)\varphi_i \in X_{ki}$. By definition of C_{ki} , $[\varepsilon, \pi_1 \cdots \pi_k]\varphi_i$ must contain an edge (γ_i, c) of C_{ki} . Hence $[\varepsilon, \pi_1 \cdots \pi_k]$ must contain an edge (γ, c) such that $(\gamma, c)\varphi_i = (\gamma_i, c) \in C_{ki}$. By definition, $(\gamma, c) \in C_k$, as required. \square

The following technical lemma is an essential step towards our goal. Assume that $M_k \cap M_{k+1} = \emptyset$ for all $k = 1, \dots, n - 2$. While reading the proof, it may be useful to draw pictures.

Lemma 3.13. *For each $k = 1, \dots, n-1$, the following hold:*

- (1) $C_k \subseteq [1, \pi_1 \dots \pi_k]$;
- (2) *for each $(\varepsilon, e) \in [1, \pi_1 \dots \pi_k] \cap [1, \pi_1 \dots \pi_{k-1}]$, there exists $(\gamma_k, c_k) \in C_k \cap [\varepsilon, \pi_1 \dots \pi_k]$.*

Proof. The proof is by induction. When $k = 1$, (1) holds by definition of C_1 and A_1 , and (2) is vacuous.

Assume that the statement is true for some $k \in \{1, \dots, n-2\}$. We show that it is valid for $k+1$. First we claim that each edge $(\varepsilon, e) \in [1, \pi_1 \dots \pi_k] \setminus M_k$ is protected from cancellation when $\pi_1 \dots \pi_k$ is multiplied by π_{k+1} on the right. Note that

$$[1, \pi_1 \dots \pi_k] \subseteq [1, \pi_1 \dots \pi_{k-1}] \cup \pi_1 \dots \pi_{k-1} [1, \pi_k] = [1, \pi_1 \dots \pi_{k-1}] \cup A_k.$$

If $(\varepsilon, e) \in [1, \pi_1 \dots \pi_{k-1}]$, then, by (2), there exists $(\gamma_k, c_k) \in C_k \cap [\varepsilon, \pi_1 \dots \pi_k]$. If $(\varepsilon, e) \in A_k$, then, since $(\varepsilon, e) \notin M_k$, we may again deduce that there exists $(\gamma_k, c_k) \in C_k \cap [\varepsilon, \pi_1 \dots \pi_k]$, this time by Lemma 3.12. Now, using $(\gamma_k, c_k) \notin A_{k+1}$ (recall: $C_k \cap A_{k+1} = \emptyset$), we may deduce, by Lemma 3.9, that $(\varepsilon, e) \in [1, \pi_1 \dots \pi_{k+1}]$. Notice that we have established the following:

$$(3.2) \quad \text{if } (\varepsilon, e) \in [1, \pi_1 \dots \pi_k] \setminus M_k, \text{ then there exists } (\gamma_k, c_k) \in C_k \cap [\varepsilon, \pi_1 \dots \pi_k].$$

We now proceed to show, as a consequence, that all elements of $A_{k+1} \setminus M_k$ are protected from cancellation as well. Indeed, take any edge $(\varepsilon, e) \in A_{k+1} \setminus M_k$; if $(\varepsilon, e) \notin [1, \pi_1 \dots \pi_{k+1}]$, then $(\varepsilon, e) \in [1, \pi_1 \dots \pi_k]$ (since $A_{k+1} \subseteq [\pi_1 \dots \pi_k, 1] \cup [1, \pi_1 \dots \pi_k \pi_{k+1}]$). But then, by the above argument, $(\varepsilon, e) \in M_k$, a contradiction. Hence, $(\varepsilon, e) \in A_{k+1} \setminus M_k$ implies $(\varepsilon, e) \in [1, \pi_1 \dots \pi_{k+1}]$. We thus have deduced (1) for $k+1$ since $C_{k+1} \subseteq A_{k+1} \setminus M_k$.

It remains to prove (2) for $k+1$. We state an intermediary result first. Take any $(\gamma_k, c_k) \in C_k$ and any $(\beta_{k+1}, b_{k+1}) \in B_{k+1}$; then $(\beta_{k+1}, b_{k+1}) \notin A_k$. As a consequence,

$$(\beta_{k+1}, b_{k+1}) \in [\gamma_k, \pi_1 \dots \pi_{k+1}].$$

Indeed, we can apply Lemma 3.8(2) with $\beta = \beta_{k+1}$, $b = b_{k+1}$, $\pi = \pi_1 \dots \pi_k$, $\sigma = \pi_{k+1}$, and $\gamma = \gamma_k$ because of the following facts: $[\gamma, \pi] \subseteq A_k$ (since $\pi = \pi_1 \dots \pi_k \in A_k$ and $(\gamma_k, c_k) \in C_k \subseteq A_k$), whence $(\beta, b) \notin [\gamma, \pi]$ but $(\beta, b) \in B_{k+1} \subseteq A_{k+1} = \pi[1, \sigma]$.

Since $B_{k+1} \cap M_{k+1} = \emptyset$, there exists, by Lemma 3.12,

$$(\gamma_{k+1}, c_{k+1}) \in C_{k+1} \cap [\beta_{k+1}, \pi_1 \dots \pi_{k+1}]$$

and so

$$(\gamma_{k+1}, c_{k+1}) \in [\beta_{k+1}, \pi_1 \dots \pi_{k+1}] \subseteq [\gamma_k, \pi_1 \dots \pi_{k+1}].$$

Altogether we have shown that

$$(3.3) \quad \text{for each } (\gamma_k, c_k) \in C_k, \text{ there exists } (\gamma_{k+1}, c_{k+1}) \in C_{k+1} \cap [\gamma_k, \pi_1 \dots \pi_{k+1}].$$

Now take any $(\varepsilon, e) \in [1, \pi_1 \dots \pi_k \pi_{k+1}] \cap [1, \pi_1 \dots \pi_k]$. Then, by (3.2), either $(\varepsilon, e) \in M_k$, or there exists $(\gamma_k, c_k) \in C_k$ such that $(\gamma_k, c_k) \in [\varepsilon, \pi_1 \dots \pi_k]$. In the latter case, since

$$[\varepsilon, \pi_1 \dots \pi_k] \subseteq [\varepsilon, \pi_1 \dots \pi_k \pi_{k+1}] \cup \pi_1 \dots \pi_k [1, \pi_{k+1}] = [\varepsilon, \pi_1 \dots \pi_{k+1}] \cup A_{k+1}$$

and $(\gamma_k, c_k) \notin A_{k+1}$, $(\gamma_k, c_k) \in [\varepsilon, \pi_1 \cdots \pi_k \pi_{k+1}]$; therefore $[\gamma_k, \pi_1 \cdots \pi_{k+1}] \subseteq [\varepsilon, \pi_1 \cdots \pi_{k+1}]$. By (3.3), there exists $(\gamma_{k+1}, c_{k+1}) \in C_{k+1}$ such that $(\gamma_{k+1}, c_{k+1}) \in [\gamma_k, \pi_1 \cdots \pi_{k+1}]$, from which we obtain

$$(\gamma_{k+1}, c_{k+1}) \in [\gamma_k, \pi_1 \cdots \pi_{k+1}] \subseteq [\varepsilon, \pi_1 \cdots \pi_{k+1}],$$

as required.

So we are left with the case

$$(\varepsilon, e) \in [1, \pi_1 \cdots \pi_{k+1}] \cap [1, \pi_1 \cdots \pi_k] \cap M_k.$$

Since $M_k = A_k \cap A_{k+1} \subseteq A_{k+1}$ and $M_k \cap M_{k+1} = \emptyset$, we may use Lemma 3.12 to conclude that there exists $(\gamma_{k+1}, c_{k+1}) \in [\varepsilon, \pi_1 \cdots \pi_{k+1}] \cap C_{k+1}$, as required. \square

We are now able to complete the proof of Theorem 3.7.

Proof. (3.1) is equivalent to

$$[\pi_1 \cdots \pi_{k-2}, \pi_1 \cdots \pi_{k-1}] \cap [\pi_1 \cdots \pi_{k-1}, \pi_1 \cdots \pi_k] \cap [\pi_1 \cdots \pi_k, \pi_1 \cdots \pi_{k+1}] = \emptyset,$$

that is,

$$A_{k-1} \cap A_k \cap A_{k+1} = \emptyset.$$

This in turn is equivalent to $M_{k-1} \cap M_k = \emptyset$ since $M_{k-1} = A_{k-1} \cap A_k$ and $M_k = A_k \cap A_{k+1}$. We may now apply Lemma 3.13 to obtain $\emptyset \neq C_{n-1} \subseteq [1, \pi_1 \cdots \pi_{n-1}]$. Since $C_{n-1} \cap A_n = \emptyset$, an application of Lemma 3.8(1) gives that $\emptyset \neq C_{n-1} \subseteq [1, \pi_1 \cdots \pi_{n-1} \pi_n]$ and hence $\pi_1 \cdots \pi_n \neq 1$, proving the claim. \square

We mention that the “unique factorization results” in relatively free profinite groups obtained by Almeida and Weil [5], [6] are immediate consequences of Theorem 3.7.

4. INVERSE MONOIDS

Our primary tool for understanding and constructing arboreous pseudovarieties of groups will be inverse monoids. This is not surprising taking into account recent research indicating that inverse monoids encode useful information about the geometry of groups (cf. [13], [27], [29], [52], [54]). In this section, we collect all the information on inverse monoids which will be needed in the sequel. We also establish a connection between the property of a pseudovariety of groups being arboreous and a structural property of relatively free profinite inverse monoids.

Recall that an *inverse monoid* [25] is a monoid I with involution $x \mapsto x^{-1}$ such that $xx^{-1}x = x$ and $xx^{-1}yy^{-1} = yy^{-1}xx^{-1}$ hold for all $x, y \in I$. In particular, any group is an inverse monoid; in fact, an inverse monoid is a group if and only if 1 is its only idempotent. We use $E(I)$ for the set of idempotents of I ; this is an idempotent, commutative submonoid of I — that is, a semilattice. It is known that x^{-1} is the unique element y such that $xyx = x$, $xyy = y$, from which it is clear that any monoid morphism of inverse monoids preserves the involution. Also, any quotient of an inverse monoid by a monoid congruence is an inverse monoid [25].

It is well known that every inverse monoid is isomorphic to a monoid of partial injective mappings on a set, with composition of partial mappings as the multiplication and inversion of injective mappings as the involution [25]. A *profinite* inverse monoid is a projective limit of finite inverse monoids, or, equivalently, a compact, totally disconnected topological inverse monoid.

There is a natural partial order on any inverse monoid defined by $x \leq y$ if $xx^{-1}y = x$, or, equivalently, $x \in E(I)y$, or, equivalently, $x \in yE(I)$ [25]. This order is compatible with both the multiplication and the involution [25] and is preserved by any morphism. In each representation of an inverse monoid as a monoid of partial injective mappings, this natural order coincides with set theoretic inclusion (the mappings being regarded as binary relations). Every inverse monoid I has a *least group congruence* σ . It is defined by $x \sigma y$ if x and y have a common lower bound with respect to the natural order. One calls I/σ the *maximal group image* of I .

If $\varphi : M \rightarrow N$ is a morphism of inverse monoids, one sets $\ker \varphi = E(N)\varphi^{-1}$. An inverse monoid I is called *E-unitary* if, for the canonical morphism $\sigma^\natural : I \rightarrow I/\sigma$, $\ker \sigma^\natural = E(I)$. An inverse monoid M is a *cover* of N if there is a surjective morphism $\varphi : M \rightarrow N$ with $\varphi|_{E(M)}$ injective (which then, in fact, is bijective). The covering theorem [30], [25] shows that every inverse monoid has an *E-unitary cover*; if the inverse monoid is finite, the *E-unitary cover* can be taken to be finite.

Since inverse monoids form a variety of monoids with involution, there is a free inverse monoid $FIM(A)$ on any set A . It is known that $FIM(A)$ is *E-unitary* with maximal group image the free group $FG(A)$ on A . The structure of this inverse monoid is quite transparent and motivates the connection with tree-like graphs. We denote the free monoid on a set X by X^* . Munn's solution to the word problem for $FIM(A)$ [32] shows that two words $u, v \in \tilde{A}^*$ map to the same element of $FIM(A)$ if and only if they map to the same element of $FG(A)$ and in $\Gamma_A(FG(A))$ they traverse the same set of (geometric) edges from 1. Thus $FIM(A)$ can be described more succinctly as the set of all pairs (X, g) with $g \in FG(A)$ and X a finite connected subgraph of $\Gamma_A(FG(A))$ containing 1 and g . The product is given by $(X, g)(Y, h) = (X \cup gY, gh)$ and the involution by $(X, g)^{-1} = (g^{-1}X, g^{-1})$. A generator a is sent to $(\{(1, a)\}, a)$, where $\{(1, a)\}$ is the subgraph of $\Gamma_A(G)$ with positive edge set $\{(1, a)\}$. The natural partial order is given by $(X, g) \leq (Y, h)$ if and only if $g = h$ and $Y \subseteq X$. The natural projection to $FG(A)$ is the projection to the second coordinate; the kernel is $E(FIM(A)) = \{(X, g) \in FIM(A) \mid g = 1\}$.

This construction was generalized by Margolis and Meakin [26]. If G is an A -generated abstract group (perhaps infinite), one defines the inverse monoid $M(G)$ to consist of all pairs (X, g) with $g \in G$ and X a finite connected subgraph in $\Gamma_A(G)$ containing 1 and g . The multiplication and involution are as above, as is the natural partial order. One has that $M(G)$ is A -generated (as above), *E-unitary* with maximal group image G , and is the freest A -generated inverse monoid with this property [26]. More precisely this means the following: if M , generated by A , is an *E-unitary* inverse monoid with $M/\sigma = G$ via a morphism $\varphi : M \rightarrow G$ respecting generators, then there is a (unique) morphism $\psi : M(G) \rightarrow M$ respecting generators. Moreover, ψ is idempotent pure (and surjective). For example, $FIM(A) = M(FG(A))$. Also, it is shown in [26] that $M(G)$ is residually finite. The word problems for $M(G)$ and G are equivalent; see, for instance, [13]. We note that $M(G)$ actually depends on the generating set A of G . We also remark that the construction is functorial in the appropriate category.

This construction was further generalized by Almeida and Weil [7]. If G is a profinite group, (topologically) generated by a profinite set A , then let $\widehat{M}(G)$ consist of all pairs (X, g) with $g \in G$ and X a connection between 1 and g in $\Gamma_A(G)$. If one defines a multiplication and involution on $\widehat{M}(G)$ as above for $M(G)$, then one gets

a profinite E -unitary inverse monoid; moreover, the natural partial order is a closed relation on $\widehat{M}(G)$. Note that the topology on $\widehat{M}(G)$ is inherited from the product topology on $\mathcal{C}_{\Gamma_A(G)}(1, 1) \times G$ where, as in section 2, $\mathcal{C}_{\Gamma_A(G)}(1, 1)$ is the space of all connected (profinite) subgraphs of $\Gamma_A(G)$ containing the vertex 1 endowed with the topology described in that section. Moreover $\widehat{M}(G)$ is the projective limit of the monoids $M(G_i)$ over all finite quotients G_i of G . The maximal group image of $\widehat{M}(G)$ is G and, as in the abstract case, $\widehat{M}(G)$ is the freest E -unitary profinite inverse monoid with maximal group image G . In particular, $\widehat{M}(\widehat{F}_{\mathbf{G}}(A))$ is the free profinite inverse monoid, denoted $\widehat{FIM}(A)$.

If G is a group, then we say that G gives rise to an E -unitary cover of the inverse monoid N , or that the inverse monoid N has an E -unitary cover over G , if there is an E -unitary cover M of N such that $G = M/\sigma$. If \mathbf{H} is a pseudovariety of groups, then the collection of all finite inverse monoids with E -unitary covers over (a group in) \mathbf{H} is a pseudovariety of inverse monoids denoted $\mathbf{SI} * \mathbf{H}$. The motivation for this name is that \mathbf{SI} is the pseudovariety of semilattices, $*$ denotes the semidirect product of pseudovarieties, and it turns out that E -unitary inverse monoids always embed in a semidirect product of their maximal group image with a semilattice [25].

It is shown in [7], [48], that $\widehat{M}(\widehat{F}_{\mathbf{H}}(A))$ is the free pro- $\mathbf{SI} * \mathbf{H}$ inverse monoid on A , which we shall denote $\widehat{FIM}_{\mathbf{H}}(A)$. If A is a finite set, we let $FIM_{\mathbf{H}}(A)$ denote the relatively free inverse monoid in the (Birkhoff) variety generated by $\mathbf{SI} * \mathbf{H}$. It is known (see [26]) that $FIM_{\mathbf{H}}(A) = M(F_{\mathbf{H}}(A))$. Also, $FIM_{\mathbf{H}}(A)$, endowed with the pro- $\mathbf{SI} * \mathbf{H}$ topology, embeds topologically in $\widehat{FIM}_{\mathbf{H}}(A)$ (A a finite set). Since the variety of all inverse monoids is generated by the class of all finite inverse monoids ($FIM(A) = M(FG(A))$ is residually finite for each A , as mentioned earlier), the fact that each finite inverse monoid has a finite E -unitary cover shows that $FIM_{\mathbf{G}}(A) = FIM(A)$ and $\widehat{FIM}_{\mathbf{G}}(A) = \widehat{FIM}(A)$.

To explain the relationship between arboreous pseudovarieties and inverse monoids, we now isolate an inverse monoid theoretic property of $FIM(A) = M(FG(A))$ which states precisely that $\Gamma_A(FG(A))$ is a tree. Observe that the σ -class of $M(G)$ corresponding to an element $g \in G$ consists of all pairs (X, g) with X a finite connection between 1 and g in $\Gamma_A(G)$. Furthermore, the natural partial order is given by reverse inclusion. Thus $\Gamma_A(G)$ is “tree-like” if and only if each σ -class of $M(G)$ has a greatest element. An inverse monoid in which every σ -class has a greatest element is called F -inverse [25]; such are necessarily E -unitary. A completely analogous argument shows the following.

Theorem 4.1. *Let G be a profinite group, generated by a profinite set A . Then $\Gamma_A(G)$ is tree-like if and only if $\widehat{M}(G)$ is an F -inverse monoid. In particular, \mathbf{H} is arboreous if and only if $\widehat{FIM}_{\mathbf{H}}(A)$ is F -inverse for every finite set A .*

The fact that, for an arborescent pseudovariety \mathbf{H} , $\widehat{FIM}_{\mathbf{H}}(A)$ is F -inverse was first observed and exploited in [48], [52].

Suppose now that G is dendral with respect to a profinite set of generators A . For convenience, we let, for $g \in G$, $\widetilde{g} = ([1, g], g) \in \widehat{M}(G)$. Note that $\widetilde{gh} \leq \widetilde{g}h$, $\widetilde{1} = 1$, and $\widetilde{g}^{-1} = \widetilde{g}^{-1}$. Thus $g \mapsto \widetilde{g}$ is a dual prehomomorphism in the sense of [25]; however, it is *not* continuous. In the case that $G = \widehat{F}_{\mathbf{H}}(A)$ for a non-trivial arboreous pseudovariety \mathbf{H} , if $w \in FG(A)$ (viewed as a reduced word), then Proposition 2.12 shows that \widetilde{w} is the element of $FIM_{\mathbf{H}}(A) = FIM(A)$ represented

by the reduced word w . It is therefore natural to think of the elements of $\widehat{M}(G)$ of the form \widetilde{g} as being reduced. We observe that, for $\pi \in G$, a factorization $\pi = \alpha\beta$ is *reduced* in the sense of Almeida and Weil [5], [6] (that is, $[1, \alpha\beta] = [1, \alpha] \cup \alpha[1, \beta]$) precisely when $\widetilde{\alpha\beta} = \widetilde{\alpha}\widetilde{\beta}$. The following factorization result can then be viewed as an interpretation of Theorem 3.7 inside the profinite inverse monoid $\widehat{M}(G)$. We shall use this in the proof of the main result (Theorem 5.14) of section 5.

Theorem 4.2. *Let G be dendral with respect to a profinite generating set A . Let $\pi_1, \dots, \pi_n \in G$, $n \geq 3$, be such that $\pi_1 \cdots \pi_n = 1$. Then there exists $k \in \{2, \dots, n-1\}$ and a reduced factorization $\pi_k = \alpha\beta$ such that, in $\widehat{M}(G)$,*

$$\widetilde{\pi_{k-1}} = \widetilde{\pi_{k-1}}\widetilde{\alpha}\widetilde{\alpha}^{-1} \text{ and } \widetilde{\pi_{k+1}} = \widetilde{\beta}^{-1}\widetilde{\beta}\widetilde{\pi_{k+1}}.$$

Proof. By Theorem 3.7, there exists $k \in \{2, \dots, n-1\}$ such that

$$[1, \pi_{k-1}] \cap [\pi_{k-1}, \pi_{k-1}\pi_k] \cap [\pi_{k-1}\pi_k, \pi_{k-1}\pi_k\pi_{k+1}] \neq \emptyset.$$

Suppose that π is a vertex in this intersection (we recommend that the reader draw pictures). Let $\alpha = \pi_{k-1}^{-1}\pi$, $\beta = \pi^{-1}\pi_{k-1}\pi_k$. Then $\pi_k = \alpha\beta$ and

$$\widetilde{\alpha\beta} = ([1, \pi_{k-1}^{-1}\pi] \cup [\pi_{k-1}^{-1}\pi, \pi_k], \pi_k) = ([1, \pi_k], \pi_k) = \widetilde{\pi_k}$$

by Proposition 2.2 (4), since $\pi_{k-1}^{-1}\pi \in [1, \pi_k]$ by choice of π . Thus the factorization $\alpha\beta$ is reduced. Since $\pi_{k-1}\alpha = \pi \in [1, \pi_{k-1}]$ we also have $\pi_{k-1}[1, \alpha] = [\pi_{k-1}, \pi_{k-1}\alpha] \subseteq [1, \pi_{k-1}]$ and $\pi_{k-1}\alpha[1, \alpha^{-1}] = [\pi_{k-1}\alpha, \pi_{k-1}] \subseteq [1, \pi_{k-1}]$. Therefore

$$\begin{aligned} \widetilde{\pi_{k-1}}\widetilde{\alpha}\widetilde{\alpha}^{-1} &= ([1, \pi_{k-1}], \pi_{k-1})([1, \alpha], \alpha)([1, \alpha^{-1}], \alpha^{-1}) \\ &= ([1, \pi_{k-1}] \cup \pi_{k-1}[1, \alpha] \cup \pi_{k-1}\alpha[1, \alpha^{-1}], \pi_{k-1}) \\ &= ([1, \pi_{k-1}], \pi_{k-1}) = \widetilde{\pi_{k-1}}. \end{aligned}$$

The remaining statement is proved analogously. □

5. CONNECTIONS WITH SEPARABILITY PROPERTIES OF FREE GROUPS

In the main theorem of this section, we shall characterize arboreous pseudovarieties in various ways. Several of the stated equivalent conditions are in terms of inverse monoids. We also establish a connection to separability properties of free groups: the property of being arboreous for a non-trivial pseudovariety of groups \mathbf{H} is shown to be equivalent to a certain separability condition for the pro- \mathbf{H} -topology on the free group $FG(A)$ which is stronger than being n -coset separable.

For this purpose, we need to recall some definitions. An *inverse automaton* \mathcal{A} over a finite alphabet A consists of a connected graph, also denoted \mathcal{A} , and a labeling of the edges by \widetilde{A} such that: the positive edges are labeled by A ; the negative edges are labeled by A^{-1} ; $a \in \widetilde{A}$ labels the edge e if and only if a^{-1} labels the edge e^{-1} ; and no two edges with the same label enter (or exit) any vertex. We observe that the labeling can be viewed as a graph immersion to a bouquet of $|A|$ circles [13], [27], [45], [50].

If $p = e_1 \dots e_n$ is a path, then the sequence $w = a_1 \dots a_n$ of corresponding labels is called the *label* of p (we say that w labels the path p). The labels then give rise to partial one-to-one functions on the vertex set $V(\mathcal{A})$ by defining, for $w \in \widetilde{A}^*$, $q \in V(\mathcal{A})$, qw to be the terminal vertex of the unique path labeled by w from q , if such a path exists, and undefined, otherwise. The collection of all partial functions

arising in this manner is an A -generated inverse monoid, denoted $I(\mathcal{A})$, called the *transition inverse monoid* of \mathcal{A} . If \mathcal{A} is finite, then so is $I(\mathcal{A})$.

Next, for a vertex i of \mathcal{A} , let $\pi_1(\mathcal{A}, i)$ be the set of all elements $w \in FG(A)$ (viewed as reduced words, including 1), which label a circuit p with $ip = \tau p = i$. Then $\pi_1(\mathcal{A}, i)$ is a subgroup of $FG(A)$, the *fundamental group of \mathcal{A} (with respect to the base point i)*; changing the base point i leads to a conjugate subgroup. We usually omit the base point and denote by $\pi_1(\mathcal{A})$ the *fundamental group of \mathcal{A}* , which is unique up to conjugation and coincides with the usual topological notion of the fundamental group of the underlying graph. If \mathcal{A} is finite, then $\pi_1(\mathcal{A})$ is finitely generated.

One says that a subset of $FG(A)$ is \mathbf{H} -closed for a pseudovariety \mathbf{H} if it is closed with respect to the pro- \mathbf{H} topology. We observe that, since conjugation is a homeomorphism, the fundamental group of an inverse automaton is \mathbf{H} -closed at one vertex if and only if it is at every vertex. The following result was proved in [29] for so-called reduced automata and, in general, by the second author in [50].

Theorem 5.1. *Let \mathbf{H} be a pseudovariety of groups and \mathcal{A} be a finite inverse automaton over an alphabet A with \mathbf{H} -closed fundamental group in $FG(A)$. Then the transition inverse monoid $I(\mathcal{A})$ is in $\mathbf{Sl} * \mathbf{H}$.*

Following [52], one says that \mathbf{H} is a *Hall pseudovariety* if the converse of Theorem 5.1 holds; that is, for any finite inverse automaton \mathcal{A} over a finite set A , if the transition inverse monoid $I(\mathcal{A})$ has an E -unitary cover over \mathbf{H} , then $\pi_1(\mathcal{A})$ is an \mathbf{H} -closed subgroup of $FG(A)$.

In light of the covering theorem for finite inverse monoids, the fact that \mathbf{G} is a Hall pseudovariety follows from Hall's famous theorem [18] stating that each finitely generated subgroup of a free group is \mathbf{G} -closed (hence the name). If \mathbf{H} is extension-closed, it is Hall by [29]; the corresponding result for arborescent pseudovarieties is [52, Theorem 9.2]. The results of [29] imply that if \mathbf{H} is a Hall pseudovariety, then the \mathbf{H} -closure of a finitely generated subgroup of a free group is again finitely generated. It is shown in [29], [57] that \mathbf{G}_{nil} is *not* Hall (although the \mathbf{G}_{nil} -closure of a finitely generated subgroup of $FG(A)$ is finitely generated). By using similar arguments to those in [29], [57] — in particular, by a careful adaptation of [57, Example 3.7] — one can prove the following.

Theorem 5.2. *Any Hall pseudovariety of nilpotent groups is a pseudovariety of p -groups.*

It is interesting to note that the property of being Hall for a pseudovariety \mathbf{H} can also be expressed in terms of a geometric property of the Cayley graphs of the finitely generated free pro- \mathbf{H} groups, a property which, for an abstract graph, reduces to being a tree. It will then be immediate that each non-trivial arboreous pseudovariety is Hall.

Theorem 5.3. *A pseudovariety \mathbf{H} of groups is Hall if and only if, for each finite alphabet A , any one of the following equivalent conditions is satisfied:*

- (1) *For each connected subgraph Δ of a quotient graph of the Cayley graph of an A -generated group in \mathbf{H} , the fundamental group of Δ is \mathbf{H} -closed in $FG(A)$.*
- (2) *For each connected subgraph Δ of the Cayley graph of an A -generated group in \mathbf{H} , the fundamental group of Δ is \mathbf{H} -closed in $FG(A)$.*

- (3) Each free generator of $\widehat{FIM}_{\mathbf{H}}(A)$ is the (unique) greatest element in its σ -class.
- (4) If e is an edge of $\Gamma_A(\widehat{F}_{\mathbf{H}}(A))$ and Λ is any connection between ιe and τe in $\Gamma_A(\widehat{F}_{\mathbf{H}}(A))$, then $e \in \Lambda$.
- (5) \mathbf{H} residually contains the free group $FG(A)$ and any element $w \in FIM(A) \subseteq \widehat{FIM}_{\mathbf{H}}(A)$ is the greatest element in its σ -class in $\widehat{FIM}_{\mathbf{H}}(A)$.

Proof. It is well known that the transition monoid of a finite inverse automaton belongs to $\mathbf{SI} * \mathbf{H}$ if and only if it is a subgraph of a quotient graph of the Cayley graph of a group in \mathbf{H} ; see [50, Proposition 5.3], for instance. Thus (1) is equivalent to \mathbf{H} being Hall.

Clearly (1) implies (2). Suppose that (2) holds and let $a \in A$ and $(\Delta, a) \in \widehat{FIM}_{\mathbf{H}}(A)$. To verify (3) we need to show that $(\{(1, a), (a, a^{-1})\}, a) \geq (\Delta, a)$, that is, the edge $(1, a) \in \Gamma_A(\widehat{F}_{\mathbf{H}}(A))$ belongs to Δ . For a word $w \in \tilde{A}^*$, let $[w]_{\mathbf{H}}$ be the image of w in $F_{\mathbf{H}}(A)$ and $(\widehat{w}, [w]_{\mathbf{H}})$ the image in $FIM_{\mathbf{H}}(A) \subseteq \widehat{FIM}_{\mathbf{H}}(A)$. Let $\{w_n\}_{n \in \mathbb{N}} \subseteq \tilde{A}^*$ be a sequence of words such that $\lim(\widehat{w}_n, [w_n]_{\mathbf{H}}) = (\Delta, a)$. Let $G \in \mathbf{H}$ be A -generated and $\psi : \widehat{FIM}_{\mathbf{H}}(A) \rightarrow M(G)$ be the canonical morphism. The sequence $(\widehat{w}_n, [w_n]_{\mathbf{H}})\psi$ is eventually constant with value, say, (Δ, a) . We need to show that the edge $(1, a)$ is in Δ .

Without loss of generality we may assume that $(\widehat{w}_n, [w_n]_{\mathbf{H}})\psi = (\Delta, a)$ for all $n \in \mathbb{N}$. That is, each w_n labels a path in $\Gamma_A(G)$ from 1 to a and uses precisely the (geometric) edges of Δ . Moreover, $[w_n]_{\mathbf{H}} \rightarrow a$ in $F_{\mathbf{H}}(A)$, whence $[w_n]_{FG(A)} \rightarrow [a]_{FG(A)}$ in the pro- \mathbf{H} -topology of $FG(A)$. By continuity of multiplication, $[w_n w_1^{-1}]_{FG(A)} \rightarrow [a w_1^{-1}]_{FG(A)}$. Let u_n and u be, respectively, the reduced forms of $w_n w_1^{-1}$ and $a w_1^{-1}$, and consider these as elements of $FG(A)$. By construction $u_n \in \pi_1(\Delta, 1)$, so by (2), $u \in \pi_1(\Delta, 1)$. If the first letter of u is a , then $(1, a) \in \Delta$ and we are done. If not, then w_1^{-1} is of the form $v_1 a^{-1} v_2$, where v_1 reduces to 1 in $FG(A)$; but w_1^{-1} labels a path in Δ from a to 1 and v_1 reads a circuit at a . Consequently $a^{-1} v_2$ labels a path in Δ from a to 1 and so $(a, a^{-1}) = (1, a)^{-1} \in \Delta$. Since Δ is a subgraph of $\Gamma_A(G)$, we may conclude that $(1, a) \in \Delta$. As G was arbitrary, it follows that $(1, a) \in \Delta$.

Clearly (3) is equivalent to (4). Suppose that (4) holds. Immediately we have that \mathbf{H} is non-trivial. We first prove that \mathbf{H} residually contains $FG(A)$ for any finite set A . Indeed, if $F_{\mathbf{H}}(A)$ were not free, then there would exist a minimal circuit p of length at least two in the Cayley graph $\Gamma_A(F_{\mathbf{H}}(A)) \subseteq \Gamma_A(\widehat{F}_{\mathbf{H}}(A))$. Let e be the first edge of this circuit; so $p = eq$. Then q^{-1} is a connection between ιe and τe not containing e , a contradiction. We now show that every reduced word w over \tilde{A} is maximal in its σ -class of $\widehat{FIM}_{\mathbf{H}}(A)$. If u is a reduced word, we use $[1, u]$ to denote the geodesic from 1 to u in $\Gamma_A(FG(A))$. Let (Δ, w) be in the σ -class of w and suppose that $[1, w] \not\subseteq \Delta$. Let e be an edge of $[1, w] \setminus \Delta$ and $\Lambda = ([1, w] \cup \Delta) \setminus \{e\}$ be the graph obtained from $[1, w] \cup \Delta$ by removing the edge e (but retaining all vertices). Then Λ is a connection between ιe and τe not containing e , a contradiction. That (5) implies the Hall property is the content of the proof of [52, Theorem 9.2]. \square

We say that \mathbf{H} enjoys *n-coset separability* if, for every finite set A and every finite collection H_1, \dots, H_n of finitely generated, \mathbf{H} -closed subgroups of $FG(A)$, the product $H_1 \cdots H_n$ is an \mathbf{H} -closed subset of $FG(A)$. The special case $n = 2$ is of particular importance and is usually referred to as *double coset separability*.

The fact that \mathbf{G} enjoys double coset separability is the fact that $FG(A)$ is double coset separable; see [16], [34], [39]. This was generalized by the second author to arborescent pseudovarieties \mathbf{H} [52, Corollary 9.4]. An unpublished example of Almeida and the second author shows that the product of two \mathbf{G}_{nil} -closed subgroups may be proper, but \mathbf{G}_{nil} -dense (showing that \mathbf{G}_{nil} does not enjoy double coset separability). That \mathbf{G} enjoys n -coset separability for each n is the striking theorem of Ribes and Zalesskiĭ [39] which led to the confirmation of the long-standing Rhodes Type II conjecture [20]. Ash's independent proof [8] of that conjecture led in turn to another proof of the n -coset separability of \mathbf{G} , made explicit by the second author [52]. A third proof of that important result, by model theoretic means, was found by Herwig and Lascar [19]. Ribes and Zalesskiĭ actually proved a more general result, namely that each extension-closed pseudovariety \mathbf{H} enjoys n -coset separability for all n [40]. One of the main results in this section will be to extend this result to all arboreous pseudovarieties.

In light of Theorems 4.1 and 5.3, the proof of [52, Corollary 9.4] goes through without change for arboreous pseudovarieties. Hence we may deduce the following.

Theorem 5.4. *Each non-trivial arboreous pseudovariety of groups is Hall and enjoys double coset separability.*

We now work towards the converse of this result. In the process, we shall come upon other equivalent conditions to being arboreous that will enable us to construct (further down the road) a large class of examples. Our main tool is again inverse monoids.

First we need the following definition [14]. Let $\varphi : M \rightarrow N$ be a relation of monoids. Then φ is a *relational morphism* if: $m\varphi \neq \emptyset$ for all $m \in M$, $1 \in 1\varphi$, and $m\varphi \cdot n\varphi \subseteq (mn)\varphi$ for all $m, n \in M$. That is, $\#\varphi$, the *graph* of φ , is a submonoid of $M \times N$ projecting onto M . If M and N are inverse monoids and $\#\varphi$ is an inverse submonoid of $M \times N$, then φ is a *relational morphism of inverse monoids*. Relational morphisms arise by considering, for some monoid R , a surjective morphism $\tau : R \rightarrow M$ and a morphism $\alpha : R \rightarrow N$ and setting $\varphi = \tau^{-1}\alpha$ (one can take $R = \#\varphi \subseteq M \times N$). The following easy lemma shows that, for our purposes, we need not distinguish between “relational morphism” and “relational morphism of inverse monoids”.

Lemma 5.5 ([50, Proposition 5.2]). *Let $\psi : S \rightarrow G$ be a relational morphism with S an inverse monoid and G a torsion group. Then ψ is a relational morphism of inverse monoids.*

Fix a finite inverse monoid S and a pseudovariety of groups \mathbf{H} . A subset $X \subseteq S$ is said to be \mathbf{H} -*pointlike* or *pointlike with respect to \mathbf{H}* if, for any relational morphism $\varphi : S \rightarrow G \in \mathbf{H}$, there exists $g \in G$ such that $X \subseteq g\varphi^{-1}$; we will often simply say that a set is *pointlike* if it is pointlike with respect to \mathbf{G} .

If $\varphi : S \rightarrow T$ is a relational morphism, one says that $X \subseteq S$ is *pointlike* for φ if there exists $t \in T$ such that $X \subseteq t\varphi^{-1}$. If $G \in \mathbf{H}$ and $\varphi : S \rightarrow G$ is such that $X \subseteq S$ is \mathbf{H} -pointlike if and only if it is pointlike for φ , then we say that φ (respectively G) *computes, or witnesses, the \mathbf{H} -pointlike sets* of S . It is well known that, for each S , there is actually a group $G \in \mathbf{H}$ and a relational morphism $\varphi : S \rightarrow G$ computing the \mathbf{H} -pointlike sets of S . Namely, for each set $X \subseteq S$ which is not \mathbf{H} -pointlike, choose a group $G_X \in \mathbf{H}$ and a relational morphism $\varphi_X : S \rightarrow G_X$ such that X

is not contained in any $g\varphi_X^{-1}$. Then take $G = \prod G_X$ and $\varphi = \prod \varphi_X$, where the products range over the (finitely many) subsets of S that are not \mathbf{H} -pointlike.

The following result is elementary.

Proposition 5.6. *Let \mathbf{H} be any pseudovariety of groups and suppose the elements of $X \subseteq S$ have a common upper bound in the natural partial order. Then X is \mathbf{H} -pointlike.*

Proof. Let $\varphi : S \rightarrow G \in \mathbf{H}$ be a relational morphism. Suppose that b is a common upper bound for X . Choose any element $(b, g) \in \#\varphi$ and, for each $x \in X$, choose $(x, g_x) \in \#\varphi$. Then $(x, g_x)(x, g_x)^{-1}(b, g) = (x, g)$ for all $x \in X$, since $x \leq b$ implies $xx^{-1}b = x$. Since each element on the left-hand side is in $\#\varphi$, so is (x, g) . Since x was an arbitrary element of X , $X \subseteq g\varphi^{-1}$. It follows that X is \mathbf{H} -pointlike. \square

The rest of this section is organized as follows. First we establish a relationship between Hall pseudovarieties enjoying double coset separability and the converse of Proposition 5.6 for \mathbf{H} -pointlike pairs. Then we show that, for such \mathbf{H} , the relatively free pro- $\mathbf{Sl} * \mathbf{H}$ inverse monoids are F -inverse which in turn implies that such pseudovarieties are arboreous. Finally we obtain that arboreous pseudovarieties enjoy n -coset separability for each n . An essential ingredient for this will be the concept of an \mathbf{H} -liftable n -tuple.

An n -tuple $(a_1, \dots, a_n) \in S^n$ is said to be \mathbf{H} -liftable [20], [52] if, for any relational morphism $\varphi : S \rightarrow G \in \mathbf{H}$, $1 \in a_1\varphi \cdot a_2\varphi \cdots a_n\varphi$. We simply call a tuple liftable if it is \mathbf{G} -liftable. \mathbf{H} -liftable 2-tuples and \mathbf{H} -pointlike pairs stand in the following relation.

Lemma 5.7 ([50, Lemma 7.23]). *Let \mathbf{H} be a pseudovariety of groups. Then $\{a, c\} \subseteq S$ is \mathbf{H} -pointlike if and only if (a, c^{-1}) is an \mathbf{H} -liftable 2-tuple.*

Throughout, we fix the following notation. Suppose that the inverse monoid S is generated by A and let $\rho : \tilde{A}^* \rightarrow FG(A)$ and $\alpha : \tilde{A}^* \rightarrow S$ be the canonical morphisms. For $m \in S$, set

$$L_m = m\alpha^{-1}.$$

If $S \in \mathbf{Sl} * \mathbf{H}$ and $(m_1, \dots, m_n) \in S^n$ is such that $1 \in L_{m_1}\rho \cdot L_{m_2}\rho \cdots L_{m_n}\rho$, then one can easily verify that (m_1, \dots, m_n) is \mathbf{H} -liftable [8], [50]. The following, which is [52, Theorem 8.2], relates the converse to n -coset separability.

Theorem 5.8. *Let \mathbf{H} be a Hall pseudovariety. Then \mathbf{H} enjoys n -coset separability if and only if, for each A -generated inverse monoid $S \in \mathbf{Sl} * \mathbf{H}$ and each \mathbf{H} -liftable n -tuple $(m_1, \dots, m_n) \in S^n$, $1 \in L_{m_1}\rho \cdot L_{m_2}\rho \cdots L_{m_n}\rho$.*

An immediate consequence of Lemma 5.7 and Theorem 5.8, combined with the observation that, for $A, B \subseteq FG(A)$, $1 \in AB$ if and only if $A \cap B^{-1} \neq \emptyset$, is the following:

Corollary 5.9. *A Hall pseudovariety \mathbf{H} enjoys double coset separability if and only if, for each A -generated inverse monoid $S \in \mathbf{Sl} * \mathbf{H}$ and each \mathbf{H} -pointlike pair $\{a, c\} \subseteq S$, $L_a\rho \cap L_c\rho \neq \emptyset$.*

This allows us to prove the next statement.

Proposition 5.10. *Let \mathbf{H} be a Hall pseudovariety enjoying double coset separability; let $S \in \mathbf{Sl} * \mathbf{H}$ and $a, b \in S$. If $\{a, b\}$ is \mathbf{H} -pointlike, then a and b have a common upper bound with respect to the natural order \leq .*

Proof. Let \mathbf{H} be a Hall pseudovariety enjoying double coset separability, let S be a finite A -generated inverse monoid in $\mathbf{Sl} * \mathbf{H}$, and suppose that $\{a, b\} \subseteq S$ is \mathbf{H} -pointlike. By Corollary 5.9, there exists $w \in FG(A)$ with $w \in L_a \rho \cap L_b \rho$. The element w can be represented as a reduced word in \tilde{A}^* , which we also denote by w . Then, at the same time, this reduced word w represents a unique element of $FIM(A)$, which we also denote by w . Let $v_a, v_b \in \tilde{A}^*$ be words in L_a, L_b , respectively, which reduce (in $FG(A)$) to w . Then, in $FIM(A)$, $v_a, v_b \leq w$ (reduced words in free inverse monoids are greatest in their σ -class); but then $a, b \leq w\alpha$ since $a = v_a\alpha$ and $b = v_b\alpha$. \square

We can now relate the conditions we have been considering to arboreous pseudovarieties.

Theorem 5.11. *Let A be a finite set and suppose that, for each A -generated inverse monoid $S \in \mathbf{Sl} * \mathbf{H}$, every \mathbf{H} -pointlike pair $\{a, b\} \subseteq S$ has an upper bound in S . Then $\widehat{FIM}_{\mathbf{H}}(A)$ is F -inverse.*

Proof. By Theorem 4.1, we need to show that $\Gamma_A(\widehat{F}_{\mathbf{H}}(A))$ is tree-like. By Proposition 2.6, it suffices to show that any two connections between vertices $g, h \in \widehat{F}_{\mathbf{H}}(A)$ contain another connection. Left translation allows us to assume that $h = 1$. Suppose that X, Y are two connections between 1 and g .

Let $S \in \mathbf{Sl} * \mathbf{H}$ be A -generated and let $\psi_S : \widehat{FIM}_{\mathbf{H}}(A) \rightarrow S$ be the canonical morphism. Let

$$C_S = \{x \in \widehat{FIM}_{\mathbf{H}}(A) \mid (X, g)\psi_S, (Y, g)\psi_S \leq x\psi_S\}.$$

First note that C_S is non-empty since a standard argument (cf. [2], [46], [47]) shows that the set $\{(X, g)\psi_S, (Y, g)\psi_S\}$ is \mathbf{H} -pointlike and hence, by assumption, has an upper bound. We observe that C_S is closed, being the union of the finitely many clopen sets $b\psi_S^{-1}$, where b is a common upper bound of $(X, g)\psi_S$ and $(Y, g)\psi_S$. Observe that the collection of the sets C_S , where S ranges over all A -generated inverse monoids in $\mathbf{Sl} * \mathbf{H}$, satisfies the finite intersection condition. Indeed, if S_1, \dots, S_n are A -generated inverse monoids in $\mathbf{Sl} * \mathbf{H}$, and T is the inverse submonoid of $S_1 \times \dots \times S_n$ generated by the “diagonal embedding” of A , then it is easy to see that $C_T \subseteq C_{S_1} \cap \dots \cap C_{S_n}$. It follows that $\bigcap C_S$ is non-empty. Since in a profinite inverse monoid $x \leq y$ if and only if this inequality holds in all finite images, it follows that any $(Z, k) \in \bigcap C_S$ is a common upper bound of (X, g) and (Y, g) . But then $k = g$ and $Z \subseteq X \cap Y$ and so Z is the desired connection. \square

We can now show the converse of Proposition 5.6 for arboreous pseudovarieties.

Theorem 5.12. *Suppose that \mathbf{H} is arboreous and let $S \in \mathbf{Sl} * \mathbf{H}$. If $X \subseteq S$ is \mathbf{H} -pointlike, then X has an upper bound.*

Proof. Let $S \in \mathbf{Sl} * \mathbf{H}$ be A -generated and $X \subseteq S$ be \mathbf{H} -pointlike. Let $\psi : \widehat{FIM}_{\mathbf{H}}(A) \rightarrow S$ and $\tau : \widehat{FIM}_{\mathbf{H}}(A) \rightarrow \widehat{F}_{\mathbf{H}}(A)$ be the canonical projections. A standard compactness argument [46] shows that there exists $g \in \widehat{F}_{\mathbf{H}}(A)$ such that $X \subseteq g\tau^{-1}\psi$. Since $([1, g], g)$ is the greatest element of $g\tau^{-1}$ it follows that $([1, g], g)\psi$ is an upper bound for X . \square

We now intend to prove that arboreous pseudovarieties enjoy n -coset separability for each n . For this purpose we introduce the following sets of pseudovarieties of

groups. For each positive integer n , let \mathcal{C}_n be the collection of all pseudovarieties of groups \mathbf{H} such that, for each finite alphabet A and for each A -generated inverse monoid $S \in \mathbf{SI} * \mathbf{H}$, $1 \in L_{m_1}\rho \cdots L_{m_n}\rho$ for each \mathbf{H} -liftable n -tuple $(m_1, \dots, m_n) \in S^n$ (or, equivalently, by Ash's theorem [8], for any inverse monoid in $\mathbf{SI} * \mathbf{H}$, the \mathbf{H} -liftable and the liftable n -tuples coincide).

We then have a descending chain $\mathcal{C}_1 \supseteq \mathcal{C}_2 \supseteq \mathcal{C}_3 \supseteq \dots$ of sets of pseudovarieties of groups, where \mathcal{C}_1 is the set of all pseudovarieties and \mathcal{C}_2 is the set of all arboreous pseudovarieties. Indeed, if $\mathbf{H} \in \mathcal{C}_2$, $S \in \mathbf{SI} * \mathbf{H}$, and $\{a, b\} \subseteq S$ is \mathbf{H} -pointlike, then (a, b^{-1}) is \mathbf{H} -liftable (Lemma 5.7) whence, by assumption, $1 \in L_a\rho L_{b^{-1}}\rho$ so that $L_a\rho \cap L_b\rho \neq \emptyset$. By the argument in the proof of Proposition 5.10, one then gets that $\widehat{\{a, b\}}$ has an upper bound, so \mathbf{H} satisfies the hypothesis of Theorem 5.11. Then $\widehat{FIM}_{\mathbf{H}}(A)$ is F -inverse so, by Theorem 4.1, \mathbf{H} is arboreous. Conversely, the trivial pseudovariety $\mathbf{1}$ is in \mathcal{C}_2 since, for a semilattice, any pair of elements (a, b) is a $\mathbf{1}$ -liftable 2-tuple, and $1 \in L_a\rho L_b\rho$ holds for all a, b . If \mathbf{H} is arboreous and non-trivial, then, by Theorem 5.4, \mathbf{H} is Hall and enjoys double coset separability so, by Theorem 5.8, $\mathbf{H} \in \mathcal{C}_2$. In particular, each non-trivial member of \mathcal{C}_2 and therefore of \mathcal{C}_n is Hall. By Theorem 5.8, all members of \mathcal{C}_n enjoy n -coset separability. In the following, we prove that $\mathcal{C}_2 = \mathcal{C}_n$ for all $n \geq 2$. The proof is inspired by [8]; see also [52].

Proposition 5.13. $\mathcal{C}_n = \mathcal{C}_2$ for each $n \geq 2$.

Proof. We show that $\mathcal{C}_{n-1} \subseteq \mathcal{C}_n$ for each $n > 2$; the claim then follows since the reverse inclusion is trivial for each n . Let $\mathbf{H} \in \mathcal{C}_{n-1}$ and $S \in \mathbf{SI} * \mathbf{H}$ be A -generated. We use $\psi : \widehat{FIM}_{\mathbf{H}}(A) \rightarrow S$ and $\tau : \widehat{FIM}_{\mathbf{H}}(A) \rightarrow \widehat{F}_{\mathbf{H}}(A)$ for the canonical continuous surjections. A standard compactness argument (cf. [2], [47]) shows that $(m_1, \dots, m_k) \in S^k$ is an \mathbf{H} -liftable k -tuple if and only if

$$1 \in m_1\psi^{-1}\tau \cdots m_k\psi^{-1}\tau.$$

Suppose that $(s_1, \dots, s_n) \in S^n$ is \mathbf{H} -liftable. Then there exist $\pi_k \in s_k\psi^{-1}\tau$ ($k = 1, \dots, n$) such that $\pi_1 \cdots \pi_n = 1$. Let $r_k = \widetilde{\pi}_k\psi$ ($k = 1, \dots, n$), where, for each $g \in \widehat{F}_{\mathbf{H}}(A)$, $\widetilde{g} = ([1, g], g) \in \widehat{FIM}_{\mathbf{H}}(A)$, as at the end of section 4. Then (r_1, \dots, r_n) is \mathbf{H} -liftable, as well. We first work with this n -tuple.

Let $k \in \{2, \dots, n-1\}$ and the reduced factorization $\pi_k = \alpha\beta$ be as per Theorem 4.2. For $i \leq k-2$ and $i \geq k+2$, let $\pi'_i = \pi_i$ and put $\pi'_{k-1} = \pi_{k-1}\alpha$, and $\pi'_{k+1} = \beta\pi_{k+1}$. Then

$$\pi'_1 \cdots \pi'_{k-1}\pi'_{k+1} \cdots \pi'_n = \pi_1 \cdots \pi_n = 1,$$

whence $(r_1, \dots, r_{k-2}, r_{k-1} \cdot \widetilde{\alpha}\psi, \widetilde{\beta}\psi \cdot r_{k+1}, r_{k+2}, \dots, r_n)$ is an \mathbf{H} -liftable $(n-1)$ -tuple. So, by assumption on \mathbf{H} , there exist $w_1, \dots, w_{k-2}, w_{k+2}, \dots, w_n, r, s \in \widetilde{A}^*$ such that

$$w_1 \cdots w_{k-2} r s w_{k+2} \cdots w_n = 1 \text{ holds in } FG(A)$$

and $w_i\psi = r_i$ for $i \neq k-1, k, k+1$, $r\psi = r_{k-1} \cdot \widetilde{\alpha}\psi$, and $s\psi = \widetilde{\beta}\psi \cdot r_{k+1}$. Let $u, w \in \widetilde{A}^*$ be such that $u\psi = \widetilde{\alpha}\psi$ and $w\psi = \widetilde{\beta}\psi$. Set $w_{k-1} = ru^{-1}$, $w_k = uw$, and $w_{k+1} = w^{-1}s$. It is easy to verify, using the conclusions of Theorem 4.2, that $w_{k-1}\psi = r_{k-1}$, $w_k\psi = r_k$, and $w_{k+1}\psi = r_{k+1}$. Moreover, in $FG(A)$ the following equalities hold:

$$w_1 \cdots w_n = w_1 \cdots w_{k-2} r u^{-1} u w w^{-1} s w_{k+2} \cdots w_n = 1.$$

To finish the proof, choose, for $i = 1, \dots, n$, words $v_i \in \tilde{A}^*$ such that $v_i\psi = s_i$. Since $s_i \leq r_i$ for all i (by construction), $(v_i v_i^{-1} w_i)\psi = s_i s_i^{-1} r_i = s_i$. Setting $t_i = v_i v_i^{-1} w_i$ for all i , we have that

$$t_1 \cdots t_n = w_1 \cdots w_n = 1$$

holds in $FG(A)$. The result follows. □

We may summarize the results of this section as follows:

Theorem 5.14. *For a non-trivial pseudovariety \mathbf{H} of groups the following conditions are equivalent:*

- (1) \mathbf{H} is arboreous;
- (2) \mathbf{H} is Hall and enjoys double coset separability;
- (3) \mathbf{H} is Hall and enjoys n -coset separability for each $n \geq 2$;
- (4) $\widehat{FIM}_{\mathbf{H}}(A)$ is F -inverse for every finite set A ;
- (5) for every inverse monoid $S \in \mathbf{SI} * \mathbf{H}$, \mathbf{H} -pointlike pairs of elements of S have an upper bound;
- (6) for every inverse monoid $S \in \mathbf{SI} * \mathbf{H}$, \mathbf{H} -pointlike subsets of S have an upper bound;
- (7) for every inverse monoid $S \in \mathbf{SI} * \mathbf{H}$, the \mathbf{H} -liftable 2-tuples are the \mathbf{G} -liftable 2-tuples;
- (8) for every inverse monoid $S \in \mathbf{SI} * \mathbf{H}$ and each $n \geq 2$, the \mathbf{H} -liftable n -tuples are the \mathbf{G} -liftable n -tuples.

6. BEHAVIOR UNDER SEMIDIRECT PRODUCT

In this section, we study the behavior of semidirect products of pseudovarieties of groups with respect to the property of being arboreous. We first prove a preliminary result showing that, to compute the liftable n -tuples (pointlike sets) of a finite inverse monoid, it suffices to find a relational morphism to a non-trivial group (say, for instance, one induced by an E -unitary cover) and then compute the liftable n -tuples (pointlike sets) for Tilson’s derived semigroup [14], [56] of the relational morphism. As a consequence, we show that the set of all non-trivial arboreous pseudovarieties forms a right ideal in the monoid of all groups pseudovarieties. The method of proof is via inverse monoids, exploiting Theorems 5.8 and 5.14.

We recall that, for pseudovarieties of groups \mathbf{V} and \mathbf{W} , their product $\mathbf{V} * \mathbf{W}$ is the pseudovariety generated by all semidirect products of members of \mathbf{V} by members of \mathbf{W} ; equivalently, it consists of all extensions of groups in \mathbf{V} by groups in \mathbf{W} . This product is associative; see [14]. The proofs are inspired by [49, Theorem 4.1], where the pointlike sets for semidirect products are described in terms of pointlike sets of the factors.

If A is a set and S is an inverse monoid, we use $B_A^1(S)$ to denote the inverse monoid obtained by adjoining an identity to the Rees matrix semigroup over S with $|A| \times |A|$ identity structure matrix. More precisely, $B_A^1(S)$ consists of the set $(A \times S \times A) \cup \{1, 0\}$ where 1 is the identity, 0 is the zero, and the non-zero products are given by

$$(a_1, s_1, a'_1)(a'_1, s_2, a_2) = (a_1, s_1 s_2, a_2);$$

the involution is given by $(a_1, s, a_2)^{-1} = (a_2, s^{-1}, a_1)$. Let $\varphi : I \rightarrow G$ be a morphism of inverse monoids with G a group. The *derived inverse monoid* of φ is the inverse submonoid D_φ of $B_G^1(I)$ consisting of 1, 0 and all elements of the form $(g, t, g \cdot t\varphi)$

(this is the derived semigroup [14] of φ with an adjoined identity, which is essentially the consolidation of the derived category of φ [56]).

Let I be an inverse monoid and G be a group; a relational morphism $\gamma : I \rightarrow G$ (respectively, the group G) *computes, or witnesses, the liftable n -tuples* of I provided that, for any $(m_1, \dots, m_n) \in I^n$, $1 \in m_1\gamma \cdots m_n\gamma$ implies that (m_1, \dots, m_n) is liftable. For each inverse monoid I , there always exists such a relational morphism (respectively, group).¹

Theorem 6.1. *Let I be an inverse monoid and $\varphi : I \rightarrow G_2$ be a surjective morphism with G_2 a non-trivial group. Suppose, furthermore, that G_1 computes the liftable n -tuples (pointlike sets) of D_φ . Then the wreath product $G_1 \wr G_2$ computes the liftable n -tuples (pointlike sets) of I .*

Proof. We just handle the case of liftable n -tuples as the argument for pointlike sets is similar (and easier). Let A be a generating set for I , which we then take as a generating set for G_2 . Let $\gamma : D_\varphi \rightarrow G_1$ compute the liftable n -tuples of D_φ . Let $W = G_1 \wr G_2$ and define $\zeta : I \rightarrow W$ by

$$i\zeta = \{(f, i\varphi) \in G_1^{G_2} \times G_2 \mid gf \in (g, i, g \cdot i\varphi)\gamma \text{ for all } g \in G_2\}.$$

To see that this is a relational morphism, first observe that, for each $g \in G_2$, $(g, 1, g \cdot 1\varphi) = (g, 1, g)$ is an idempotent of D_φ and hence γ -relates to 1. Thus 1 ζ -relates to the identity of W . Suppose that $(f_1, t_1\varphi) \in t_1\zeta$, $(f_2, t_2\varphi) \in t_2\zeta$. Then

$$(f_1, t_1\varphi)(f_2, t_2\varphi) = (f_1 \cdot {}^{t_1}\varphi f_2, (t_1 t_2)\varphi).$$

So if $g \in G_2$, then

$$\begin{aligned} g(f_1 \cdot {}^{t_1}\varphi f_2) &= gf_1 \cdot (g \cdot t_1\varphi)f_2 \in (g, t_1, g \cdot t_1\varphi)\gamma \cdot (g \cdot t_1\varphi, t_2, g(t_1\varphi \cdot t_2\varphi))\gamma \\ &\subseteq (g, t_1 t_2, g \cdot (t_1 t_2)\varphi)\gamma, \end{aligned}$$

as desired. It follows that ζ is a relational morphism.

Let $(r_1, \dots, r_n) \in I^n$ be such that $(\bar{1}, 1) \in r_1\zeta \cdots r_n\zeta$, where $\bar{1} : G_1 \rightarrow G_2$ is the constant function sending all elements to 1 (so $(\bar{1}, 1)$ is the identity of W). Let $g_j = r_j\varphi$ and suppose that $(f_j, g_j) \in r_j\zeta$ are such that

$$(\bar{1}, 1) = (f_1, g_1) \cdots (f_n, g_n) = (f_1 {}^{g_1}\varphi f_2 \cdots {}^{g_1 \cdots g_{n-1}}\varphi f_n, g_1 \cdots g_n).$$

Evaluating the first component at 1, we see that

$$1 = (1)f_1 \cdot (g_1)f_2 \cdots (g_1 \cdots g_{n-1})f_n.$$

But the definition of ζ implies

$$(g_1 \cdots g_{j-1})f_j \in (g_1 \cdots g_{j-1}, r_j, g_1 \cdots g_j)\gamma$$

for each j (where, for $j = 1$, the empty product $g_1 \cdots g_0$ is interpreted as 1). So, by choice of γ ,

$$((1, r_1, g_1), \dots, (g_1 \cdots g_{n-1}, r_n, g_1 \cdots g_n)) \in D_\varphi^n$$

is a liftable n -tuple.

For convenience, in the following we shall identify an element of A with its image in I . We now aim to show that D_φ is generated by $B = G_2 \times (A \cup \{1\})$ via the map $\lambda : (g, a) \mapsto (g, a, g \cdot a\varphi)$.² It is convenient to identify B^{-1} with the set of all pairs

¹The argument is completely analogous to the argument showing that there is always a relational morphism computing the pointlike sets.

²This essentially is the well-known fact that the derived semigroupoid (or category) of φ is generated by the graph $\Gamma_A(G_2)$.

$(g \cdot a\varphi, a^{-1})$ with $a \in A \cup \{1\}$ (here 1^{-1} is distinct from 1 as a symbol); observe that $(g, 1)$ and $(g, 1^{-1})$ both have image $(g, 1, g)$ in D_φ . To prove our claim, note that any element $(g, i, g \cdot i\varphi)$ can be written as a product

$$\begin{aligned} & (g, a_1)\lambda \cdot (g \cdot a_1\varphi, a_2)\lambda \cdots (g \cdot (a_1 \cdots a_{k-1})\varphi, a_k)\lambda \\ &= (g, a_1, g \cdot a_1\varphi)(g \cdot a_1\varphi, a_2, g \cdot (a_1 a_2)\varphi) \cdots (g \cdot (a_1 \cdots a_{k-1})\varphi, a_k, g \cdot (a_1 \cdots a_k)\varphi) \\ &= (g, a_1 \cdots a_k, g \cdot (a_1 \cdots a_k)\varphi) \\ &= (g, i, g \cdot i\varphi), \end{aligned}$$

where $i = a_1 \cdots a_k$ ($a_j \in \tilde{A} \cup \{1\}$). Since G_2 is non-trivial, $0 = (1, 1)\lambda \cdot (g, 1)\lambda$ for any $g \in G_2, g \neq 1$.

Consider the map $\eta : \tilde{B} \rightarrow \tilde{A}^*$ given by

$$(g, a) \mapsto a, (g \cdot a\varphi, a^{-1}) \mapsto a^{-1}, (g, 1^{\pm 1}) \mapsto 1.$$

This map preserves inverses, so there is an induced morphism $\eta : \tilde{B}^* \rightarrow \tilde{A}^*$ of monoids with involution (abusing notation). A simple calculation shows that if $(h_1, a_1)\lambda \cdots (h_k, a_k)\lambda = (h_1, i, h_1 \cdot i\varphi)$, then $a_1 \cdots a_k = i$ in I (where 1^{-1} is interpreted as 1); that is, the value of $((h_1, a_1) \cdots (h_k, a_k))\eta$ in I is i .

Since

$$((1, r_1, g_1), \dots, (g_1 \cdots g_{n-1}, r_n, g_1 \cdots g_n)) \in D_\varphi^n$$

is a liftable n -tuple, there exist $u_1, \dots, u_n \in \tilde{B}^*$ such that $u_j\lambda = (g_1 \cdots g_{j-1}, r_j, g_1 \cdots g_j)$ for each j and the word $u_1 \cdots u_n$ reduces to 1 in $FG(B)$; observe that none of the words u_j is empty. Since u_j does not map to 0 in D_φ , our above observation shows that the value of $u_j\eta$ in I is r_j . Since $u_1 \cdots u_n$ reduces to 1 in $FG(B)$ and η is a morphism of free monoids with involution, $u_1\eta \cdots u_n\eta$ also reduces to 1. We conclude that (r_1, \dots, r_n) is liftable. \square

Suppose that $\zeta : I \rightarrow G$ computes the liftable n -tuples (pointlike sets) and $\varphi : I \rightarrow S$ is a surjective morphism of inverse monoids. By a standard argument $\varphi^{-1}\zeta : S \rightarrow G$ computes the liftable n -tuples (pointlike sets) of S [37, Lemma 1.4]. Therefore Theorem 6.1 essentially says that in order to find a relational morphism witnessing the liftable n -tuples of an A -generated inverse monoid S , we need only: (i) find an A -generated group G which gives rise to an E -unitary cover of S and (ii) find a group H which calculates the liftable n -tuples of D_φ , where $\varphi : M(G) \rightarrow G$ is the canonical morphism. Note that in this case D_φ has a very transparent structure. Besides 0 and 1, its elements are of the form (g, B, h) , where $g, h \in G$ and B is a connected subgraph of the Cayley graph of G containing g, h . The non-zero products of such elements are given by

$$(g, B, h)(h, C, k) = (g, B \cup C, k),$$

and the involution is defined by $(g, B, h)^{-1} = (h, B, g)$.

We may now prove the main theorem of this section.

Theorem 6.2. *The set of all non-trivial arboreous pseudovarieties is a right ideal in the monoid of all pseudovarieties of groups.*

Proof. Let \mathbf{V} be a non-trivial arboreous pseudovariety groups and \mathbf{W} be an arbitrary pseudovariety of groups. To show that $\mathbf{V} * \mathbf{W}$ is arboreous, it suffices to show, by Theorem 5.14, that, for each $S \in \mathbf{SI} * (\mathbf{V} * \mathbf{W})$, there is a group $G \in \mathbf{V} * \mathbf{W}$ computing the liftable n -tuples of S (for some integer $n \geq 2$). Let $I \in \mathbf{SI} * (\mathbf{V} * \mathbf{W})$ be an E -unitary cover of S with a surjective, idempotent separating morphism

$\tau : I \rightarrow S$ and a surjective, idempotent pure morphism $\psi : I \rightarrow G \in \mathbf{V} * \mathbf{W}$. Then there is a surjective morphism $\rho : G \rightarrow G_2$ with $G_2 \in \mathbf{W}$ and $N = \ker \rho \in \mathbf{V}$ (N might be trivial). Let I be A -generated. Then S, G_1 , and G_2 are A -generated in a natural way. Let $\varphi = \psi\rho$. If G_2 is trivial, then $S \in \mathbf{S1} * \mathbf{V}$ and we are done since, by assumption, there is a group in $\mathbf{V} \subseteq \mathbf{V} * \mathbf{W}$ which computes the liftable n -tuples of S .

Suppose that G_2 is non-trivial. We begin by showing that $D_\varphi \in \mathbf{S1} * \mathbf{V}$. First choose a set-theoretic section $\beta : G_2 \rightarrow G$ of ρ , that is, a mapping β is such that $\beta\rho = \text{id}_{G_2}$. Since \mathbf{V} is non-trivial, it contains a cyclic group $\mathbb{Z}/p\mathbb{Z}$ for some prime p and therefore contains \mathbf{Ab}_p . Let H be the free \mathbf{Ab}_p -group generated by G_2 (written additively). Note that H is a direct sum of $|G_2|$ copies of $\mathbb{Z}/p\mathbb{Z}$ and $H \in \mathbf{V}$. Define a relational morphism $\alpha : D_\varphi \rightarrow H \times N$ by

$$0 \mapsto H \times N, 1 \mapsto (0, 1), \text{ and } (g, i, g \cdot i\varphi) \mapsto (g - g \cdot i\varphi, g\beta \cdot i\psi \cdot ((g \cdot i\varphi)\beta)^{-1}).$$

This is clearly well defined. To check that it is a relational morphism it suffices to perform the calculation

$$\begin{aligned} & (g, i, g \cdot i\varphi)\alpha \cdot (g \cdot i\varphi, i', g \cdot i\varphi i'\varphi)\alpha \\ &= (g - g \cdot i\varphi, g\beta \cdot i\psi \cdot ((g \cdot i\varphi)\beta)^{-1})(g \cdot i\varphi - g \cdot i\varphi i'\varphi, (g \cdot i\varphi)\beta \cdot i'\psi \cdot ((g \cdot i\varphi i'\varphi)\beta)^{-1}) \\ &= (g - g \cdot (ii')\varphi, g\beta \cdot (ii')\psi \cdot ((g \cdot (ii')\varphi)\beta)^{-1}) \\ &= (g, ii', g \cdot (ii')\varphi)\alpha. \end{aligned}$$

Suppose now that $(g, i, g \cdot i\varphi)\alpha$ is idempotent. Then $g - g \cdot i\varphi = 0$, whence $g = g \cdot i\varphi$. Also $1 = g\beta \cdot i\psi \cdot ((g \cdot i\varphi)\beta)^{-1} = g\beta \cdot i\psi \cdot (g\beta)^{-1}$ whence $i\psi = 1$. Since ψ is idempotent pure, i is an idempotent. It follows that the graph $\#\alpha$ of α is an E -unitary cover of D_φ over the group $H \times N \in \mathbf{V}$ and so D_φ is in $\mathbf{S1} * \mathbf{V}$.

Since \mathbf{V} is arboreous, there exists, by Theorem 5.14, a group $G_1 \in \mathbf{V}$ computing the liftable n -tuples of D_φ . Then the wreath product $G_1 \wr G_2$ belongs to $\mathbf{V} * \mathbf{W}$ and computes the liftable n -tuples of I by Theorem 6.1 and therefore of S (by [37, Lemma 1.4]). \square

In particular, since \mathbf{G}_p is extension-closed and therefore arboreous by [17], any pseudovariety \mathbf{H} of groups closed under co-extensions by p -groups for some p is arboreous; such pseudovarieties will appear again in section 9. In the next section, we give an even weaker condition for a pseudovariety to be arboreous.

It should also be noted that the set of all non-trivial arboreous pseudovarieties is by no means a left ideal. For instance, if \mathbf{V} is a pseudovariety satisfying a non-trivial group identity and \mathbf{W} is any pseudovariety such that $\mathbf{V} * \mathbf{W}$ is strictly bigger than \mathbf{W} , then $\widehat{F_{\mathbf{V} * \mathbf{W}}}$ cannot be arboreous. Indeed, there exists a finite alphabet A and $\pi \in \widehat{F_{\mathbf{V} * \mathbf{W}}}(A)$ such that $\pi \neq 1$, but π maps to 1 under the canonical morphism $\widehat{F_{\mathbf{V} * \mathbf{W}}}(A) \rightarrow \widehat{F_{\mathbf{W}}}(A)$. Write $\pi = \pi(A)$ to indicate the “variables” occurring in π . Moreover there exists a reduced word $w = w(z_1, \dots, z_n) \in FG(z_1, \dots, z_n)$ such that \mathbf{V} satisfies the identity $w = 1$. Now choose $2n$ pairwise disjoint finite alphabets $A_1, B_1, \dots, A_n, B_n$ of size $|A|$, set $X = A_1 \cup B_1 \cup \dots \cup A_n \cup B_n$ and form $\pi(A_1), \pi(B_1), \dots, \pi(A_n), \pi(B_n) \in \widehat{F_{\mathbf{V} * \mathbf{W}}}(X)$. Let

$$\nu = w(\pi(A_1)\pi(B_1), \dots, \pi(A_n)\pi(B_n)) \in \widehat{F_{\mathbf{V} * \mathbf{W}}}(X)$$

be obtained from w by substituting each z_i with $\pi(A_i)\pi(B_i)$. By construction, $\nu = 1$, whence $\widehat{F_{\mathbf{V} * \mathbf{W}}}(X)$ cannot be dendral with respect to X by Theorem 3.7.

7. A SUFFICIENT CONDITION

Here we show that a pseudovariety \mathbf{H} is arboreous if, for each $G \in \mathbf{H}$, there is a cyclic group C such that the wreath product $C \wr G$ is in \mathbf{H} , or, equivalently, for each $G \in \mathbf{H}$ there is a prime p such that \mathbf{H} contains all co-extensions of G by groups in \mathbf{Ab}_p (see Proposition 7.6).

We begin with a key graph theoretic lemma. Let Γ be an oriented graph with positive edge set E ; for a path w in Γ , let $c(w)$ be the *content* of w , that is, the set of all positive edges e such that e or e^{-1} occurs in w . For a positive edge $e \in E$ let

$$w(e) = \text{the sum of exponents of } e \text{ in } w.$$

That is, $w(e)$ is the number of signed traversals — forward minus backward — of e in w . Moreover for a prime p , let

$$w(e)_p = w(e) \pmod{p},$$

that is, $w(e)_p$ is the image of $w(e)$ under the natural morphism $\mathbb{Z} \rightarrow \mathbb{Z}/p\mathbb{Z}$. Then $w(e)_p$ denotes, modulo p , the number of signed traversals of e in w .

Lemma 7.1. *Let Γ be an oriented graph and p a prime. Then, for each path w in Γ , there exists a path u in Γ , coterminial with w , such that u uses only edges of w and, for each edge e used by u ,*

$$0 \neq u(e)_p = w(e)_p.$$

Proof. If $\iota w = \tau w$, then we may choose u to be the empty path. So consider two distinct vertices x and y of Γ ; we shall prove that our claim holds for all paths w with $\iota w = x$ and $\tau w = y$. Formally, the proof will be by induction on $(|c(w)|, |w|)$, where $|w|$ is the length of the path w and $\mathbb{N} \times \mathbb{N}$ is ordered lexicographically. Suppose first that w is a path in which each edge e is traversed in only one direction and no edge is traversed a multiple of p times. Then w itself meets the requirements, so we may take $u = w$. The claim therefore holds for all paths of this kind and such paths shall form the basis for the induction. We now consider several cases; it is recommended that the reader draw pictures.

Let w be a path with $\iota w = x$, $\tau w = y$ and suppose that the claim holds for each path w' with $\iota w' = x$ and $\tau w' = y$ such that either w' uses a smaller number of edges than w , or w' uses the same number of edges as w but is of shorter length than w . Suppose first that there is an edge e used by w in both directions. Then $w = seqe^{-1}t$ for some paths s, q, t . If no edge of q is used elsewhere in this path, then let $w' = st$; since $\tau s = \iota e = \tau e^{-1} = \iota t$, this is actually a path and clearly is coterminial with w . Moreover, each edge traversed by w' is traversed by w the same number of times. Since w' is obviously shorter than w , we apply the induction hypothesis and are done.

Now assume again that an edge e in w is used in both directions, but whenever we have a pattern $\cdots e q e^{-1} \cdots$ in w , some edge used in q is also used outside $e q e^{-1}$. In this case, it suffices to consider the following two possibilities: either (i) $w = meqfre^{-1}sf^{-1}t$ or (ii) $w = meqfre^{-1}sft$ for some edge f and some paths m, q, r, s, t . The case $w = mfqerfse^{-1}t$ is treated dually to (ii); all other possibilities are obtained by renaming the edges in (i) or (ii). In case (i), consider $w' = msrqt$. As above, it is easy to check that this actually forms a path coterminial with w . Again, all edges of $c(w')$ are traversed by w the same number of times. Since w' is shorter than w , the induction hypothesis applies. Likewise, in case (ii),

consider $w' = msfrqft$. Again, w' is a path coterminal with w , each edge in $c(w')$ is traversed the same number of times in w , and w' is shorter than w , whence the induction hypothesis applies. This exhausts all possibilities in which there is an edge used in both directions.

Thus we may assume that there is no edge in w which is used in both directions. If there is also no edge used precisely kp times for some integer k (in one direction), then we have arrived at the basis for the induction. So let us, as a last case, assume that there is an edge e used in w precisely kp times (and in one direction only). That is, $w = mec_1ec_2 \cdots ec_{kp-1}et$ for some edge e and paths $m, c_1, \dots, c_{kp-1}, t$ and e is not used in these other paths. Consider the path $w' = mc_1^{-1}c_2c_1^{-1}c_3c_1^{-1} \cdots c_1^{-1}c_{kp-1}c_1^{-1}t$. It is easily checked that w' is a path coterminal with w . It is immediate that the edge e is no longer used in w' , that is, $|c(w')| < |c(w)|$. Moreover, w' consists of the paths $m, t, c_2, \dots, c_{kp-1}$ and c_1^{-1} and the latter is traversed $kp - 1 = -1$ times modulo p . It follows that, for each edge f occurring in c_1 , the contribution of c_1 to $w(f)_p$ is the same as the contribution of the $kp - 1$ occurrences of c_1^{-1} to $w'(f)_p$. Since, for any edge f' , the other segments $m, c_2, \dots, c_{kp-1}, t$ yield the same contribution to $w(f')_p$ as they do to $w'(f')_p$, we may conclude that, for each edge f used in w' , $w'(f)_p = w(f)_p$. We therefore may apply the induction hypothesis to w' and the claim is proved. \square

Now fix a prime p , a finite set A , and an A -generated group G . Let $E = E(\Gamma_A(G)) = G \times A$ be the set of positive edges in the Cayley graph $\Gamma_A(G)$ of G . Then $F_{\mathbf{Ab}_p}(E)$ is an $|E|$ -fold direct sum of cyclic groups $\mathbb{Z}/p\mathbb{Z}$. The left action of G on E extends uniquely to an action by automorphisms on $F_{\mathbf{Ab}_p}(E)$ via $g(\sum c_e \cdot e) = \sum c_e \cdot ge$. Let $G^{\mathbf{Ab}_p}$ be the following A -generated group:

$$G^{\mathbf{Ab}_p} = \langle ((1, a), a) \mid a \in A \rangle \leq F_{\mathbf{Ab}_p}(E) \rtimes G$$

which is a co-extension of G by an Abelian group of exponent p (the semidirect product is with respect to the aforementioned action of G on $F_{\mathbf{Ab}_p}(E)$). The assignment $G \mapsto G^{\mathbf{Ab}_p}$ is a co-extension of the type considered in [15], [38], and [31]; in fact the approach of this subsection follows the program espoused in [31]. Also note that, if G is relatively free on A in \mathbf{V} , then $G^{\mathbf{Ab}_p}$ is the A -generated free object in $\mathbf{Ab}_p * \mathbf{V}$ (cf. [1]). In fact, this is a special case of the universal property of $G^{\mathbf{Ab}_p}$ established in [15], [38].

For a word $w \in \tilde{A}^*$, let $[w]_G$ be the value of w in G and \hat{w} the unique path in $\Gamma_A(G)$ starting at 1, labeled by w . It is easy to see that, with the above notation, the image of w in $G^{\mathbf{Ab}_p}$ under the morphism $\tilde{A}^* \rightarrow G^{\mathbf{Ab}_p}$ determined by $a \mapsto ((1, a), a)$ is given by

$$\left(\sum_{e \in E} \hat{w}(e)_p \cdot e, [w]_G \right).$$

Also, under the canonical morphism $\tilde{A}^* \rightarrow M(G)$, w maps to $(\widetilde{c(w)}, [w]_G)$ where $\widetilde{c(w)}$ is the subgraph of $\Gamma_A(G)$ with positive edges $c(w)$. For a path \hat{w} in $\Gamma_A(G)$, let $c_p(\hat{w})$ be the p -content of \hat{w} , that is, the set of all (positive) edges e for which $\hat{w}(e)_p \neq 0$; clearly $c_p(\hat{w}) \subseteq c(\hat{w})$ for each path \hat{w} . An immediate consequence of Lemma 7.1 is the following.

Corollary 7.2. *Let p, A, G be as above and $w \in \tilde{A}^*$. Then there exists a word $u \in \tilde{A}^*$ such that $[u]_G = [w]_G$ and $c(\hat{u}) = c_p(\hat{u}) \subseteq c_p(\hat{w})$.*

We are now able to formulate and prove the following important result.

Theorem 7.3. *Let G be an A -generated group, p a prime and $\psi : M(G) \rightarrow G^{\mathbf{Ab}_p}$ be the relational morphism induced by sending generator to generator. Then ψ computes the \mathbf{G} -pointlike sets of $M(G)$.*

Proof. According to Theorem 5.14, it suffices to show that any elements $(X_1, g_1), \dots, (X_n, g_n) \in M(G)$ which relate to a point of $G^{\mathbf{Ab}_p}$ via ψ have a common upper bound in $M(G)$. Suppose that the elements $(X_1, g_1), \dots, (X_n, g_n)$ relate to a point via ψ . Then, first of all, $g_1 = g_2 = \dots = g_n = g$ for some $g \in G$. Moreover, there are words $w_1, \dots, w_n \in \tilde{A}^*$ with w_i mapping to (X_i, g) in $M(G)$ for each i (in particular, $[w_i]_G = g$, for each i), all of which map to the same element in $G^{\mathbf{Ab}_p}$. Now the image of w_i in $G^{\mathbf{Ab}_p}$ is just

$$\left(\sum \hat{w}_i(e)_p \cdot e, g \right),$$

so it follows that, for each $e \in E(\Gamma_A(G))$, $\hat{w}_1(e)_p = \hat{w}_i(e)_p$ for all i . In particular, $c_p(\hat{w}_1) = c_p(\hat{w}_i)$ for all i . By Corollary 7.2, there exists a word $w \in \tilde{A}^*$ with $[w]_G = g$ and

$$c(\hat{w}) = c_p(\hat{w}) \subseteq c_p(\hat{w}_1) = c_p(\hat{w}_i) \subseteq c(\hat{w}_i) \subseteq c(\widetilde{\hat{w}_i}) = X_i$$

for all i . Consequently, $(c(\hat{w}), g)$, the image of w in $M(G)$, is a common upper bound of $(X_1, g), \dots, (X_n, g)$, as required. \square

As a consequence, we may deduce the following.

Corollary 7.4. *Suppose that S is an A -generated inverse monoid and G is an A -generated group such that S has an E -unitary cover over G . Then, for any prime p , the natural relational morphism $\varphi : S \rightarrow G^{\mathbf{Ab}_p}$ of A -generated inverse monoids computes the \mathbf{G} -pointlike sets of S .*

Proof. Let $\eta : M(G) \rightarrow S$ be the canonical surjective morphism mentioned in section 4; then $\varphi = \eta^{-1}\psi$, where $\psi : M(G) \rightarrow G^{\mathbf{Ab}_p}$ is the relational morphism of Theorem 7.3. The claim then follows from Theorem 7.3 and [37, Lemma 1.4]. \square

We have therefore proved the main result of this section.

Theorem 7.5. *Suppose that \mathbf{H} is a pseudovariety of groups such that, for each $G \in \mathbf{H}$, there exists a prime p with $G^{\mathbf{Ab}_p} \in \mathbf{H}$ (for some set A which generates G). Then \mathbf{H} is arboreous.*

The following proposition reformulates the above condition in several ways.

Proposition 7.6. *For a pseudovariety of groups \mathbf{H} the following are equivalent:*

- (1) *for each relatively free group $G \in \mathbf{H}$, there is a prime p such that $G^{\mathbf{Ab}_p} \in \mathbf{H}$ (for some set A freely generating G);*
- (2) *for each group $G \in \mathbf{H}$, there is a prime p such that $G^{\mathbf{Ab}_p} \in \mathbf{H}$ (for each set A generating G);*
- (3) *for each $G \in \mathbf{H}$ there is a prime p such that \mathbf{H} contains each co-extension of G by a group in \mathbf{Ab}_p ;*
- (4) *for each $G \in \mathbf{H}$, there is a non-trivial cyclic group C such that $C \wr G \in \mathbf{H}$.*

Proof. Each cyclic group C has a cyclic subgroup of prime order, isomorphic with $\mathbb{Z}/p\mathbb{Z}$ for some prime p . So $C \wr G \in \mathbf{H}$ implies $\mathbb{Z}/p\mathbb{Z} \wr G \in \mathbf{H}$. Moreover, any co-extension of G by a member $H = (\mathbb{Z}/p\mathbb{Z})^n$ of \mathbf{Ab}_p can be embedded in $(\mathbb{Z}/p\mathbb{Z})^n \wr G$ [33], [14], which in turn can be embedded in $(\mathbb{Z}/p\mathbb{Z} \wr G)^n$. So (4) implies (3) implies (2) implies (1).

Let $G \in \mathbf{H}$ be a non-trivial, A -generated group and let $C = \mathbb{Z}/n\mathbb{Z}$ be any cyclic group, which we write additively. For $g \in G$, denote by \bar{g} the function

$$\bar{g} : G \rightarrow C, g \mapsto 1, h \mapsto 0 \text{ if } h \neq g.$$

Then $C \wr G$ is generated by $\{(\bar{g}, 1), (0, a) \mid g \in G, a \in A\}$. Thus $C \wr G$ has a generating set of size $|G| + |A|$, independently of $|C|$. Let F be the $(|G| + |A|)$ -generated free object in $\langle G \rangle \subseteq \mathbf{H}$. Then there exists a prime p such that $F^{\mathbf{Ab}_p} \in \mathbf{H}$. But $F^{\mathbf{Ab}_p}$ is relatively free on $|G| + |A|$ generators in $\mathbf{Ab}_p * \langle G \rangle$, so $\mathbb{Z}/p\mathbb{Z} \wr G$ is a morpic image of $F^{\mathbf{Ab}_p}$ proving that (1) implies (4). \square

For instance, the following pseudovariety of solvable groups is arboreous:

$$(7.1) \quad \mathbf{H} = \dots \mathbf{Ab}_{11} * \mathbf{Ab}_7 * \mathbf{Ab}_5 * \mathbf{Ab}_3 * \mathbf{Ab}_2.$$

One may also alternate the semidirect product with the join operation to obtain examples which cannot be described as iterated semidirect products. Notice that each cyclic group in the above pseudovariety has square free order.

It seems natural to ask if the converse of Theorem 7.5 holds; we formulate this as a problem.

Problem 7.7. Suppose that \mathbf{H} is arboreous. Does this imply that, for each $G \in \mathbf{H}$, there exists a prime p with $G^{\mathbf{Ab}_p} \in \mathbf{H}$?

The authors have an answer to this only in very special cases. Observe that if \mathbf{H} is arboreous and $G \in \mathbf{H}$ is A -generated, then there must exist an A -generated co-extension H of G computing the pointlike pairs (sets) of $M(G)$. If \mathbf{H} is such that the morphism $H \rightarrow G$ can always be taken to have *nilpotent kernel*, then the answer is “yes”. From this it follows, for example, that, for a pseudovariety of p -groups, the conditions of Proposition 7.6 are also necessary for being arboreous; that is, for any prime p , the only arboreous pseudovariety of p -groups is the pseudovariety \mathbf{G}_p of *all* p -groups.

For a group G and a subset X of G , let $\langle X \rangle^G$ denote the normal subgroup generated by X . Another approach to Problem 7.7 would be via the following.

Problem 7.8. Let G be a profinite group which is dendral with respect to a finite set B and let $b \in B$. Is it true that $b \notin \overline{\langle B \setminus \{b\} \rangle^G}$?

A positive answer to Problem 7.8 would imply a positive answer to Problem 7.7. Note that Theorem 3.7 implies $b \notin \langle B \setminus \{b\} \rangle^G$.

8. CONNECTIONS WITH FINITE MONOID THEORY

In this section, we give a monoid theoretic characterization of arboreous pseudovarieties. It turns out that arboreous pseudovarieties are precisely the pseudovarieties of groups \mathbf{H} which satisfy the equation $\mathbf{J} * \mathbf{H} = \mathbf{J} \textcircled{m} \mathbf{H}$. Further applications to finite monoid theory are then drawn.

For undefined notions from monoid and semigroup theory, such as *Green’s relations*, *regular elements*, etc., the reader is referred to the books [1], [12], [14], [21],

[22], [25]. However, we shall explain the concepts of the *Schützenberger graph* and the *Schützenberger representation* (for more details see [55], [50]).

Let A be an alphabet and S an A -generated monoid (perhaps infinite). For $w \in A^*$, the image of w in S is denoted by $[w]_S$. Let R be an \mathcal{R} -class of S . The (right) *Schützenberger graph of R with respect to A* , denoted $\Gamma_A(S, R)$, is the labeled graph with vertex set R and edge set

$$\{(m, a) \in R \times A \mid m[a]_S \in R\}.$$

The incidence functions are given in the obvious way: $\iota(m, a) = m$ and $\tau(m, a) = m[a]_S$. As in the case of the Cayley graph of a group, a is the label of the edge (m, a) . Note that $\Gamma_A(S, R)$ is *strongly connected*: for any $m, n \in R$, there is a word $w \in A^*$ such that $m[w]_S = n$. Indeed, $\Gamma_A(S, R)$ is a strongly connected component of the (right) Cayley graph of S with respect to A (defined as in the group case), and all strongly connected components are of this form; see [50], [31] for more on monoid Cayley graphs and Schützenberger graphs.

Each $a \in A$ gives rise to a partial function $R \rightarrow R$, namely $m \mapsto m[a]_S$ provided that $m[a]_S \in R$ (or equivalently, (m, a) is an edge of $\Gamma_A(S, R)$) and the monoid generated by these partial functions is the *transition monoid* of $\Gamma_A(S, R)$, denoted $M(\Gamma_A(S, R))$. The canonical morphism $A^* \rightarrow M(\Gamma_A(S, R))$ factors through S , thus leading to a representation $S \rightarrow M(\Gamma_A(S, R))$ of S by partial transformations of R , which is called the *Schützenberger representation* of S with respect to R . This representation is independent of the generating set A (because each $s \in S$ is mapped to the partial transformation of R , $m \mapsto ms$ if $ms \in R$).

If we choose an \mathcal{R} -class R' in the \mathcal{D} -class of R , then there are $l, k \in S$ such that $r \mapsto lr$ and $s \mapsto ks$ are mutually inverse bijections between R and R' . Moreover, these bijections induce (label-preserving) isomorphisms of the corresponding Schützenberger graphs in a natural way. This implies that the Schützenberger graph and the Schützenberger representation only depend on the \mathcal{D} -class of R (and is essentially the classical Schützenberger representation discussed in [12]). Let \mathfrak{R} be a set of \mathcal{R} -classes of S and $\Gamma_A(S, \mathfrak{R})$ be the disjoint union of the graphs $\{\Gamma_A(S, R) \mid R \in \mathfrak{R}\}$. Then the transition monoid $M(\Gamma_A(S, \mathfrak{R}))$ is a subdirect product of the monoids $M(\Gamma_A(S, R))$, $R \in \mathfrak{R}$, and we have a representation

$$S \rightarrow M(\Gamma_A(S, \mathfrak{R}))$$

of S by partial transformations of $\bigcup \mathfrak{R}$. In [50], for the case where \mathfrak{R} consists of all \mathcal{R} -classes of S , the monoid $M(\Gamma_A(S, \mathfrak{R}))$ is denoted $\oplus Sch(S)$.

If S is an inverse monoid, generated by A as an inverse monoid, then $\Gamma_{\tilde{A}}(S, R)$ is naturally an oriented graph. Indeed, it is simple to verify [55], [50] that if (m, a) is an edge of $\Gamma_{\tilde{A}}(S, R)$ with $a \in \tilde{A}$, then $(m[a]_S, a^{-1})$ is an edge, as well. Moreover, $\iota(m[a]_S, a^{-1}) = m[a]_S$ and $\tau(m[a]_S, a^{-1}) = m$. Thus taking the set of all edges with label in A as the set of positive edges and defining $(m, a)^{-1} = (m[a]_S, a^{-1})$ for $a \in \tilde{A}$, we obtain an oriented graph $\Gamma_A(S, R)$ which is called the *Schützenberger graph of R with respect to (the inverse monoid generating set) A* . This graph is a connected inverse automaton (as defined in section 5). If S is a group, then there is only one \mathcal{R} -class and its Schützenberger graph is the Cayley graph $\Gamma_A(S)$.

The transition inverse monoid of $\Gamma_A(S, R)$ is denoted $I(\Gamma_A(S, R))$. Since

$$I(\Gamma_A(S, R)) = M(\Gamma_{\tilde{A}}(S, R))$$

(by definition), we have once again arrived at the Schützenberger representation of S with respect to R . The Schützenberger representation

$$S \rightarrow I(\Gamma_A(S, \mathfrak{R})) \leq \prod_{R \in \mathfrak{R}} I(\Gamma_A(S, R))$$

is well known to be faithful provided that \mathfrak{R} contains an \mathcal{R} -class from each \mathcal{D} -class of S [12].

Once again, all monoids and semigroups are taken to be finite. Let N and M be monoids and \mathbf{V} be a pseudovariety of semigroups. Then N is a \mathbf{V} -co-extension of M if there exists a surjective morphism $\varphi : N \rightarrow M$ such that, for each idempotent $e \in M$, the semigroup $e\varphi^{-1}$ is in \mathbf{V} . For pseudovarieties \mathbf{V}, \mathbf{W} of semigroups, the Mal'cev product $\mathbf{V} \circledast \mathbf{W}$ is the pseudovariety generated by all \mathbf{V} -co-extensions of members of \mathbf{W} .

8.1. The pseudovariety equation $\mathbf{J} \circledast \mathbf{H} = \mathbf{J} * \mathbf{H}$. A monoid is a *block group* if every regular element has precisely one inverse. The pseudovariety of all block groups is denoted by \mathbf{BG} . An important pseudovariety equation [20], [35] states that

$$\mathbf{BG} = \mathbf{J} * \mathbf{G} = \mathbf{J} \circledast \mathbf{G}.$$

Here \mathbf{J} denotes the pseudovariety of finite \mathcal{J} -trivial monoids and $\mathbf{J} * \mathbf{G}$ is the pseudovariety generated by all semidirect products of members of \mathbf{J} by members of \mathbf{G} . The reader should consult [20] for a history and for the various connections between the ideas we have been considering and finite monoid theory. The main result in this section (Theorem 8.3) shows that arboreous pseudovarieties are precisely those pseudovarieties of groups \mathbf{H} which are solutions to the pseudovariety equation $\mathbf{J} * \mathbf{H} = \mathbf{J} \circledast \mathbf{H}$. This finishes the work begun in [50], [52] to answer a question first posed (explicitly for certain pseudovarieties, and implicitly in general) in [35]. The importance of the pseudovariety $\mathbf{J} * \mathbf{H}$ in formal language theory is indicated in [51]: a language is recognized by a monoid in $\mathbf{J} * \mathbf{H}$ if and only if it is in the Boolean polynomial closure of the languages recognized by the groups in \mathbf{H} .

Theorem 8.1. *For each arboreous pseudovariety \mathbf{H} of groups,*

$$\mathbf{J} * \mathbf{H} = \mathbf{J} \circledast \mathbf{H}.$$

Proof. For \mathbf{H} trivial, this is clear; so assume \mathbf{H} is non-trivial. By Theorem 5.14, \mathbf{H} is Hall and enjoys double coset separability. It was pointed out in [52] that the second proof of [50, Theorem 7.24], which shows that $\mathbf{J} * \mathbf{H} = \mathbf{J} \circledast \mathbf{H}$ for extension closed pseudovarieties of groups \mathbf{H} , uses only that \mathbf{H} is Hall and enjoys double coset separability (cf. [53]). □

We now aim to prove the converse of Theorem 8.1. Let \mathbf{N}_2 be the pseudovariety of all null semigroups (recall that a semigroup is null if the product of any two elements is zero). This pseudovariety is defined by the identities $xy = yx = x^2$. Note that \mathbf{N}_2 is a subpseudovariety of \mathbf{N} , the pseudovariety of nilpotent semigroups. Also recall that $\mathbf{J} = \mathbf{N} \circledast \mathbf{SI}$ [14]. It is well known (cf. [50]) that

$$(8.1) \quad \mathbf{N} \circledast (\mathbf{SI} * \mathbf{H}) = \mathbf{N} \circledast (\mathbf{SI} \circledast \mathbf{H}) \subseteq \mathbf{J} \circledast \mathbf{H}$$

for any pseudovariety of groups \mathbf{H} . The next lemma is folklore (e.g. [11], Lemma 5.1 in combination with Lemma 6.3).

Lemma 8.2. *Let N be an \mathbf{N} -co-extension of M via the surjective morphism $\varphi : N \rightarrow M$; then φ is bijective on regular elements.*

We shall also need the following criterion for membership in $\mathbf{J} * \mathbf{H}$ [48, Theorem 8.2]: a block group N belongs to $\mathbf{J} * \mathbf{H}$ if and only if all \mathbf{H} -pointlike pairs of regular elements $\{a, b\}$ satisfy the equation

$$(8.2) \quad aa^{-1}bb^{-1} = ab^{-1}.$$

We are now prepared to prove the main result of this section.

Theorem 8.3. *A pseudovariety of groups \mathbf{H} is arboreous if and only if*

$$\mathbf{J} * \mathbf{H} = \mathbf{J} \textcircled{m} \mathbf{H}.$$

Proof. We just need to show that the above pseudovariety equation fails if \mathbf{H} is not arboreous. If \mathbf{H} is not arboreous (and hence \mathbf{H} non-trivial), then there must be an X -generated group $G \in \mathbf{H}$ such that \mathbf{H} does not compute the \mathbf{G} -pointlike pairs of $M(G)$. Indeed, by Theorem 5.14, there exists an inverse monoid $I \in \mathbf{SI} * \mathbf{H}$, generated by a set X , such that \mathbf{H} does not compute the \mathbf{G} -pointlike pairs of I . If $G \in \mathbf{H}$ is X -generated and gives rise to an E -unitary cover of I , then it follows, as in section 7, that \mathbf{H} does not compute the \mathbf{G} -pointlike pairs of $M(G)$. This means that there are elements $(A, g), (B, g) \in M(G)$ forming an \mathbf{H} -pointlike pair, but not having a common upper bound. The latter translates into saying that $A \cap B$ contains no connection between 1 and g ; in other words, 1 and g belong to different connected components of $A \cap B$.

The \mathcal{R} -classes of $M(G)$ correspond to the connected subgraphs Δ of $\Gamma_X(G)$ containing 1, and the Schützenberger graph of the \mathcal{R} -class R_Δ of $M(G)$ (corresponding to Δ) coincides with the graph Δ itself; thus we do not distinguish between $\Gamma_X(M(G), R_\Delta)$ and Δ (especially in our notation). Let \mathfrak{R} be the collection of all \mathcal{R} -classes of $M(G)$. Then $M(G)$ can be viewed as $I(\Gamma_X(M(G), \mathfrak{R}))$. Note that the vertex set Q of $\Gamma_X(M(G), \mathfrak{R})$ is, in fact, $M(G)$ itself. However, we shall call it Q to distinguish it from the inverse monoid.

Now we construct an automaton \mathcal{N} over \tilde{X} (which is no longer an inverse automaton) as follows. The vertex set is again Q . We retain all edges from $\Gamma_X(M(G), \mathfrak{R})$, both those labeled by X and those labeled by X^{-1} . For each vertex h in the connected component of $A \cap B$ containing g , and, for each edge (h, x) ($x \in \tilde{X}$) in $B \setminus A$, we add a (*non-reversible*) edge labeled by x from (A, h) to $(B, h[x]_G)$. (Note that the edges (h, x) play a role similar to the edges from the sets B_{ki} in section 3.2.) Let N be the transition monoid of \mathcal{N} . We make the following observations:

- (1) for any $w \in \tilde{X}^*$ and $q \in Q$, if qw is defined in $\Gamma_X(M(G), \mathfrak{R})$, then qw is also defined in \mathcal{N} and the value is unchanged;
- (2) if $w \in \tilde{X}^*$ and $q \in Q$ are such that qw is not defined in $\Gamma_X(M(G), \mathfrak{R})$ but is defined in \mathcal{N} , then $q = (A, h)$ and $(A, h)w = (B, h[w]_G)$;
- (3) for $w_1, w_2 \in \tilde{X}^*$ and $q \in Q$, if $[w_1]_N = [w_2]_N$ and qw_1 is defined in $\Gamma_X(M(G), \mathfrak{R})$, then qw_2 is defined in $\Gamma_X(M(G), \mathfrak{R})$, as well.

The first two observations follow by construction of \mathcal{N} ; the third observation is a consequence of the first two and the fact that, in $\Gamma_X(M(G), \mathfrak{R})$, transitions never change the first coordinate of a state.

It follows that the assignment $[w]_N \mapsto [w]_{M(G)}$ ($w \in \tilde{X}^*$) is a well-defined, surjective morphism $\varphi : N \rightarrow M(G)$ (and is, in fact, just a restriction of partial

functions). Moreover $1\varphi^{-1} = \{1\}$. We now show that, for each idempotent $e \in M(G)$, $e\varphi^{-1} \in \mathbf{N}_2$. For $e = 1$, this is trivial; so assume that $e \neq 1$. Note that if $w \in \tilde{X}^*$ maps to e , then $[w]_G = 1$. Suppose that w_1, w_2 map to e . We need to show that $[w_1w_2]_N = [w_2w_1]_N = [w_1^2]_N$. For any vertex q in Q , qw_1 is defined in $\Gamma_X(M(G), \mathfrak{R})$ if and only if qw_2 is, in which case $qw_1 = qw_2 = q$ in $\Gamma_X(M(G), \mathfrak{R})$ (and hence in \mathcal{N}); and so $qw_1^2 = qw_1w_2 = qw_2w_1 = q$ in \mathcal{N} .

Suppose that qw_1 is undefined in $\Gamma_X(M(G), \mathfrak{R})$, but is defined in \mathcal{N} . Then $q = (A, h)$ and $qw_1 = (B, h[w_1]_G) = (B, h)$ since $[w_1]_G = 1$. Thus $h \in A \cap B$. We have two cases. First suppose that the path labeled by w_1 in $\Gamma_X(G)$ from h is not contained in B . Then, since $[w_1]_{M(G)} = [w_2]_{M(G)}$ (and so the paths labeled by w_1 and w_2 from any vertex of $\Gamma_X(G)$ are coterminal and traverse the same geometric edges), the same is true for the path labeled by w_2 in $\Gamma_X(G)$ from h . Hence $(B, h)w_1$ and $(B, h)w_2$ are undefined; therefore qw_1w_2 and qw_1w_1 are not defined in \mathcal{N} . By the same reasoning, qw_2w_1 is also undefined in \mathcal{N} in this case.

For our second case, assume that the path labeled by w_1 from h in $\Gamma_X(G)$ is contained in B , whence the same is true for the path labeled by w_2 from h . First observe that h is in the connected component of $A \cap B$ containing g . Indeed, w_1 labels a path in $B \subseteq \Gamma_X(G)$ from h . As the path labeled by w_1 in \mathcal{N} from (A, h) crosses over from A to B and all new edges in \mathcal{N} are added to vertices of the form $q' = (A, h')$ with h' in the component of $A \cap B$ containing g , w_1 must travel, before it crosses the new edge, inside $A \cap B$ to some vertex of this component; hence h belongs to this component.

It follows that $w_1 = wxu$, where $x \in \tilde{X}$ labels the edge crossing over to B , and wxu is the corresponding factorization. Let f be the edge $(h[w]_G, x)$ of $\Gamma_X(G)$; by construction, $f \in B \setminus A$. As mentioned above, $(A, h)w_1 = (B, h)$ in \mathcal{N} , so

$$(A, h)w_1w_1 = (B, h)w_1 = (B, h) = (A, h)w_1w_2$$

(since $[w_2]_G = 1$ and the path labeled by w_2 in $\Gamma_X(G)$ from h is contained in B).

Now consider w_2 . Since it is equal to w_1 in $M(G)$, it must use the edge f in at least one direction. Since $f \in B \setminus A$, the path labeled by w_2 from h in $\Gamma_X(G)$ must at some point use an edge of $B \setminus A$; let $f' = (h', z)$ be the first such edge (in $\Gamma_X(G)$) and $w_2 = szt$ be the corresponding factorization. Since w_2 labels a path in B from h , s labels a path in $A \cap B$ from h to $h' = h[s]_G$. Hence h' is in the component of $A \cap B$ containing g (since h is). Thus, by construction of \mathcal{N} , there is a transition labeled by z from (A, h') to $(B, h'[z]_G)$. It follows that

$$(A, h)w_2 = (B, h'[z]_G)t = (B, h)$$

as $[w_2]_G = 1$ and w_2 labels a path from h in B . Thus $(A, h)w_2w_1 = (B, h)$, as desired. This handles all cases. We conclude that $\varphi : N \rightarrow M(G)$ is an \mathbf{N}_2 -co-extension. It follows that $N \in \mathbf{J} \textcircled{m} \mathbf{H}$ by (8.1).

Recall that if m is an element of a finite monoid, then m^ω denotes its (unique) idempotent power. We claim that there must be an \mathbf{H} -pointlike pair of regular elements $\{a, b\}$ of N , lifting (A, g) and (B, g) , respectively. Indeed, by a standard result [37, Lemma 1.4], there are lifts \bar{a} and \bar{b} of (A, g) and (B, g) which form an \mathbf{H} -pointlike pair. Choose lifts a', b' of $(A, g)^{-1}$, $(B, g)^{-1}$, respectively. Then the elements $a = \bar{a}(a'\bar{a})^\omega$ and $b = \bar{b}(b'\bar{b})^\omega$ form an \mathbf{H} -pointlike pair and are (the unique) regular lifts of (A, g) and (B, g) , respectively.

Let $u, v \in \tilde{X}^*$ map to (A, g) and (B, g) , respectively. Then $[u]_N[u^{-1}u]_N^\omega$ and $[v]_N[v^{-1}v]_N^\omega$ are regular lifts of (A, g) and (B, g) , respectively, and so must coincide

with a and b , respectively, considered above by Lemma 8.2. Also, since b^{-1} and $[v^{-1}]_N[vv^{-1}]_N^\omega$ are regular lifts of $(B, g)^{-1}$, they must coincide.

In the following, we consider the elements of N as partial transformations of Q in the obvious way (N being the transition monoid of \mathcal{N} with state set Q). We show that $ab^{-1} \neq aa^{-1}bb^{-1}$ by showing that these elements induce distinct partial transformations of Q .

By construction, $(A, 1)ab^{-1} = (B, 1)$. Indeed, $(A, 1)u = (A, g)$ in $\Gamma_X(M(G), \mathfrak{R})$, hence $(A, 1)a = (A, g)$ in \mathcal{N} . Now, since v^{-1} labels a path in $\Gamma_X(G)$ from g to 1 , contained in B , we have, by construction, that

$$(A, g)v^{-1} = (B, 1) \text{ in } \mathcal{N}.$$

Indeed, as 1 is not in the same component of $A \cap B$ as g , the path from g in $\Gamma_X(G)$ labeled by v^{-1} (which is contained in B) must eventually leave this latter component, say via an edge $f = (h, x)$. Let $v^{-1} = sxt$ be the corresponding factorization. Then the path in \mathcal{N} labeled by v^{-1} from (A, g) crosses over the (new) edge, labeled by x , with initial vertex $(A, g[s]_G) = (A, h)$ and terminal vertex $(B, h[x]_G)$, and continues (with label t) to $(B, 1)$. So

$$(A, g)b^{-1} = (A, g)[v^{-1}]_N[vv^{-1}]_N^\omega = (B, 1)[v^{-1}]_N^\omega = (B, 1),$$

the last equality holding because $(B, 1)[v^{-1}]_{M(G)} = (B, 1)$. Hence $(A, 1)ab^{-1} = (B, 1)$, as claimed.

However, $(A, 1)aa^{-1}bb^{-1}$ is undefined. Indeed, $(A, 1)aa^{-1} = (A, 1)$, since $(A, 1)uu^{-1} = (A, 1)$. We now show that $(A, 1)v$ is not defined in \mathcal{N} ; it will then follow that $(A, 1)bb^{-1}$ is not defined. Indeed, v labels a path in $\Gamma_X(G)$ from 1 to g in B ; as g is not in the same component of $A \cap B$ as 1 , v cannot label a path from 1 to g in A . Since no new edges were added in A to the component of $A \cap B$ containing 1 when forming \mathcal{N} from $\Gamma_X(M(G), \mathfrak{R})$, we see that v does not label a path from $(A, 1)$ in \mathcal{N} . It follows that $ab^{-1} \neq aa^{-1}bb^{-1}$ showing (8.2) fails, and so $N \in \mathbf{J} \textcircled{m} \mathbf{H} \setminus \mathbf{J} * \mathbf{H}$. \square

8.2. Further applications. First we shall assign to each block group $N \in \mathbf{J} \textcircled{m} \mathbf{H}$ an inverse monoid $I(N) \in \mathbf{SI} * \mathbf{H}$ as follows. Suppose that N is A -generated and let R be a regular \mathcal{R} -class. It is well known (e.g. [20], [35], [50]) that the transition monoid $M(\Gamma_A(N, R))$ is a monoid of partial injective transformations of R . Likewise, if \mathfrak{R} is a set of regular \mathcal{R} -classes, then $M(\Gamma_A(N, \mathfrak{R}))$ is a monoid of partial injective transformations of $\bigcup \mathfrak{R}$. Suppose that \mathfrak{R} contains (precisely) one \mathcal{R} -class from each regular \mathcal{D} -class. Let

$$\tau : N \rightarrow M(\Gamma_A(N, \mathfrak{R})), \quad n \mapsto \tau_n$$

be the Schützenberger representation of N (as introduced earlier) and let $I(N)$ be the inverse monoid of all partial injective mappings on $\bigcup \mathfrak{R}$ generated by $\tau(N) = \{\tau_n \mid n \in N\}$. Note that if N is generated by A (as a monoid), then $I(N)$ is also generated by A (but as an inverse monoid) via the map $a \mapsto \tau_a, a^{-1} \mapsto \tau_a^{-1}$, and τ is a morphism from N into $I(N)$. Now suppose that $\varphi : A^* \rightarrow N$ and $\alpha : \tilde{A}^* \rightarrow I(N)$ are the canonical surjections; then $\tau = \varphi^{-1}\alpha$ (viewing $A^* \subseteq \tilde{A}^*$ in the natural way). The following proposition is a variant on [48, Proposition 11.5].

Proposition 8.4. *Let \mathbf{H} be a non-trivial, arboreous pseudovariety of groups and $N \in \mathbf{J} \textcircled{m} \mathbf{H}$. Let $m_1, m_2 \in N$ be regular elements. Suppose that $\psi : I(N) \rightarrow H \in \mathbf{H}$ computes \mathbf{H} -pointlike pairs, where H is A -generated and ψ is induced by sending*

generator to generator. Then $\{m_1, m_2\}$ is \mathbf{H} -pointlike if and only if $\tau_{m_1}\psi \cap \tau_{m_2}\psi \neq \emptyset$, if and only if $\{\tau_{m_1}, \tau_{m_2}\}$ is \mathbf{H} -pointlike.

Proof. Suppose that $\{m_1, m_2\}$ is \mathbf{H} -pointlike, then $\tau_{m_1}\psi \cap \tau_{m_2}\psi \neq \emptyset$ by definition. Suppose conversely that $\tau_{m_1}\psi \cap \tau_{m_2}\psi \neq \emptyset$. Let $w_1, w_2 \in \tilde{A}^*$ be such that $[w_1]_H = [w_2]_H$ and $w_i\alpha = \tau_{m_i}$, $i = 1, 2$. By Corollary 5.9, it follows that there exist w'_1, w'_2 with $w'_i\alpha = w_i\alpha$, $i = 1, 2$, and such that $w'_1\rho = w'_2\rho = g$ where, once again, $\rho : \tilde{A}^* \rightarrow FG(A)$ is the canonical surjection. Then [50, Corollary 7.16] implies that $g \in \overline{m_i\varphi^{-1}}$, $i = 1, 2$, where we view $m_i\varphi^{-1} \subseteq FG(A)$ and where the closure is with respect to the pro- \mathbf{H} topology. It then follows immediately from the definition of \mathbf{H} -pointlike sets and the pro- \mathbf{H} topology that $\{m_1, m_2\}$ is \mathbf{H} -pointlike. \square

Our first corollary is to answer a question posed by J.-E. Pin at several conferences. He has asked for a method which, given a monoid $N \in \mathbf{J} \circledast \mathbf{G}$, would explicitly give a finite group G whose size and structure are tightly determined by N with the property that N divides a semidirect product $J \rtimes G$ with $J \in \mathbf{J}$. In the following, for a regular \mathcal{D} -class D of the block group N , denote by G_D some maximal subgroup of D and by E_D the set of idempotents of D .

Corollary 8.5. *Let N be an A -generated block group and p be any prime. Let G be any A -generated group which gives rise to an E -unitary cover of $I(N)$. Then N divides a semidirect product $J \rtimes G^{\mathbf{Ab}_p}$ with $J \in \mathbf{J}$. The size of the group G can be chosen to be at most*

$$\prod_{D \in \text{Reg}(N/\mathcal{D})} (|G_D| \cdot |E_D|)!$$

Proof. For convenience we suppose that $A \subseteq N$. The generators of $I(N)$ are of the form τ_a with $a \in A$. Consider the Schützenberger representation $\tau : N \rightarrow M(\Gamma_A(N, \mathfrak{R}))$, where \mathfrak{R} contains one \mathcal{R} -class from each regular \mathcal{D} -class. Let D be a regular \mathcal{D} -class of N and $R_D \in \mathfrak{R}$ with $R_D \subseteq D$. Extend $\tau_a|_{R_D}$ to a permutation $\hat{\tau}_a^{R_D}$ of R_D and put $\hat{\tau}_a = \bigcup \hat{\tau}_a^{R_D}$. Each \mathcal{R} -class R_D is invariant under each such permutation. Let G be the group generated by all $\hat{\tau}_a$. Then the size of G is at most $\prod |R_D|!$, the product being taken over all regular \mathcal{D} -classes, and $|R_D| = |G_D| \cdot |E_D|$. Moreover, G gives rise to an E -unitary cover of $I(N)$, so $G^{\mathbf{Ab}_p}$ computes the pointlike pairs of $I(N)$ by Corollary 7.4. Hence, by Proposition 8.4, $G^{\mathbf{Ab}_p}$ computes the pointlike pairs of regular elements for N . But the proof that $\mathbf{J} * \mathbf{G} = \mathbf{J} \circledast \mathbf{G}$ (cf. [20], [48], [53]) then shows that N divides a desired semidirect product of the form $J \rtimes G^{\mathbf{Ab}_p}$. \square

A similar application is the following, showing that $\mathbf{J} * \mathbf{H}$ and $\mathbf{J} \circledast \mathbf{H}$ are not all that far apart.

Corollary 8.6. *For each pseudovariety \mathbf{H} of groups,*

$$\mathbf{J} \circledast \mathbf{H} \subseteq \bigcap_{p \text{ prime}} \mathbf{J} * \mathbf{Ab}_p * \mathbf{H}.$$

Proof. Let $N \in \mathbf{J} \circledast \mathbf{H}$. Then, as mentioned above, $I(N) \in \mathbf{Sl} * \mathbf{H}$, whence $I(N)$ has an E -unitary cover over an A -generated group $H \in \mathbf{H}$. It follows from Corollary 8.5 that N divides a semidirect product of the form $J \rtimes H^{\mathbf{Ab}_p}$, with $J \in \mathbf{J}$, for any prime p . \square

We observe that the above inclusion is strict whenever \mathbf{H} satisfies a non-trivial group identity. Indeed, if \mathbf{H} is the trivial pseudovariety, then the 6-element Brandt monoid B_2^1 belongs to the right-hand side, but not the left-hand side. In general, the Cayley graph of $F_{\mathbf{H}}(A)$, for some finite set A , will have a minimal circuit at 1 of length at least two. Let a be the label of the first edge of the circuit and u be the label of the remainder (so \mathbf{H} satisfies $a = u^{-1}$). Let \mathcal{A} be the inverse automaton obtained by starting with the subgraph consisting of this circuit. Since $F_{\mathbf{H}}(A)$ is residually in \mathbf{H} , \mathcal{A} can be embedded in the Cayley graph of an A -generated group $G \in \mathbf{H}$. Observe that in $G^{\mathbf{Ab}_p}$ (for any prime p), $a \neq u^{-1}$ since $G^{\mathbf{Ab}_p}$ computes the pointlike pairs for $M(G)$ and the elements of $M(G)$ corresponding to a and u^{-1} have no common upper bound (au being the label of a minimal circuit). It follows that if \mathcal{B} is the inverse automaton obtained from \mathcal{A} by formally separating the terminal vertices of the path read from 1 by a and u^{-1} in \mathcal{A} , then \mathcal{B} embeds in $\Gamma_A(G^{\mathbf{Ab}_p})$. Hence the transition inverse monoid I of \mathcal{B} is in $\mathbf{Sl} * \mathbf{Ab}_p * \mathbf{H}$ for all primes p , but $I \notin \mathbf{J} \circledast \mathbf{H}$ (as an inverse monoid belongs to $\mathbf{J} \circledast \mathbf{H}$ if and only if it belongs to $\mathbf{Sl} * \mathbf{H}$ [48], [50]).

The next application generalizes some results of [36], [40] (see [48] for a detailed argument); the proof goes through without change using n -coset separability and the Hall property.

Theorem 8.7. *Let \mathbf{H} be arboreous. Then the pro- \mathbf{H} closed rational subsets of $FG(A)$ are precisely the finite unions of sets of the form $gH_1 \cdots H_n$, where H_1, \dots, H_n are pro- \mathbf{H} closed, finitely generated subgroups.*

Finally, we quote some pseudovariety equations whose proofs hold without change in the context of the various classes of pseudovarieties which we have considered. The first two of these equations follow from [5], [6]; constructive proofs have been given by the first author [9], [10] in the case that \mathbf{H} satisfies the conditions of Proposition 7.6. Recall that \mathbf{ZE} is the pseudovariety of monoids whose idempotents are central, \mathbf{ACom} is the pseudovariety of aperiodic, commutative monoids, $\overline{\mathbf{H}}$ is the pseudovariety of all monoids whose subgroups belong to \mathbf{H} , \mathbf{DH} is the pseudovariety of monoids whose regular \mathcal{D} -classes are groups in \mathbf{H} , and \mathbf{ECom} is the pseudovariety of monoids whose idempotents commute.

Theorem 8.8. *For each arboreous pseudovariety \mathbf{H} ,*

$$\mathbf{ACom} \vee \mathbf{H} = \mathbf{ZE} \cap \overline{\mathbf{H}}$$

and

$$(\mathbf{J} \cap \mathbf{ECom}) \vee \mathbf{H} = \mathbf{DH} \cap \mathbf{ECom}.$$

The following results were stated in [50] for extension closed pseudovarieties of groups, but the proofs go through for any arboreous pseudovariety. Recall that \mathbf{R} denotes the pseudovariety of \mathbf{R} -trivial monoids. See [2] for the definition of vertex rank.

Theorem 8.9. *Let \mathbf{V} be a pseudovariety of monoids and \mathbf{H} an arboreous pseudovariety of groups. Then*

$$(\mathbf{V} \cap \mathbf{R}) \circledast \mathbf{H} = (\mathbf{V} \circledast \mathbf{G}) \cap (\mathbf{R} \circledast \mathbf{H}).$$

If, in addition, \mathbf{V} has finite vertex rank, then

$$(\mathbf{V} * \mathbf{H}) \cap (\mathbf{R} * \mathbf{H}) = (\mathbf{V} * \mathbf{G}) \cap (\mathbf{R} * \mathbf{H}).$$

In particular,

$$\mathbf{Sl} * \mathbf{H} = \mathbf{ECom} \cap (\mathbf{R} * \mathbf{H}).$$

We also mention that Corollary 5.4 and Theorem 7.1 of [51] hold in this context. As a final result, we point out that the proofs of [48, Theorem 11.3, Theorem 11.4] work for any arboreous pseudovariety, whence we have the following result (see [3], [4] for undefined terminology).

Theorem 8.10. *Each arboreous pseudovariety \mathbf{H} is weakly κ -reducible.*

Moreover, establishing the hyperdecidability [2] of an arboreous pseudovariety reduces to finding an algorithm to compute the pro- \mathbf{H} closure of any finitely generated subgroup of a free group of finite rank.

9. APPENDIX: TREE-LIKE GRAPHS VERSUS PROFINITE TREES

Here we discuss the difference between the concept of a tree-like graph and the homological notion of a profinite tree. The arguments in this section are basically known — we collect them for convenience of the reader. In particular, we shall see that there are many non-trivial examples of relatively free profinite groups whose Cayley graphs are tree-like, but are *not* profinite trees in the homological sense. In this section, we shall give a graph by its vertex set V and its set of positive edges E . Throughout, the term “ring” will stand for *ring with unity*; morphisms among rings are assumed to respect the unity, and we assume that $0 \neq 1$. For a ring R and a set X , let RX denote the free R -module on X ; this may be viewed as the set of all finite formal sums $\sum r_x x$ with $x \in X$, $r_x \in R$ endowed with the obvious R -module structure. For a profinite ring R and a profinite set X , let \widehat{RX} be the free profinite R -module on the profinite set X (see [6], [42]); note that RX is always a dense submodule of \widehat{RX} .

Let $\Gamma = (E, V)$ be an abstract connected graph. It is well known that Γ is a tree if and only if, for any ring R , the unique morphism $\partial : RE \rightarrow RV$ determined by the map $e \mapsto \tau e - \iota e$ ($e \in E$) is injective. This motivates the following definition, first introduced by Gildenhuys and Ribes (see [17], [59], [6], [43]). Given a profinite ring R , a connected profinite graph $\Gamma = (E, V)$ is an *R -tree* if the unique morphism $\partial : \widehat{RE} \rightarrow \widehat{RV}$ determined by the map $e \mapsto \tau e - \iota e$ ($e \in E$) is injective. In contrast to the situation for abstract graphs, the latter definition *does* depend on the base ring R . That is, a profinite graph Γ may be an R -tree for some ring R , but not an S -tree for another ring S . The following, however, is obvious.

Lemma 9.1. *If S is a profinite subring of a profinite ring R and Γ is an R -tree, then Γ is also an S -tree.*

An analogous result (whose proof we omit) holds for quotient rings.

Lemma 9.2. *Suppose that $\varphi : R \rightarrow S$ is a surjective morphism of profinite rings and Γ is an R -tree; then Γ is also an S -tree.*

Corollary 9.3. *If Γ is an R -tree for some profinite ring R , then Γ is a $\mathbb{Z}/p\mathbb{Z}$ -tree for some prime p .*

Proof. Suppose that Γ is an R -tree; let $R' = \overline{\langle 1 \rangle}$ be the closed subgroup of $(R, +)$ generated by 1. Then R' is a subring of R and a pro-cyclic group, and $R' \neq 0$. Consequently, there is a prime p such that $\mathbb{Z}/p\mathbb{Z}$ is a quotient of R' (as an Abelian

group as well as a ring). By Lemma 9.1, Γ is an R' -tree and, by Lemma 9.2, Γ is also a $\mathbb{Z}/p\mathbb{Z}$ -tree. \square

We denote the ring, as well as the group, of p -adic integers by \mathbb{Z}_p . The next result is [59, Lemma 1.16].

Lemma 9.4. *A profinite graph Γ is a $\mathbb{Z}/p\mathbb{Z}$ -tree if and only if it is a \mathbb{Z}_p -tree.*

Altogether we have:

Corollary 9.5. *For a profinite graph Γ the following are equivalent:*

- (1) Γ is an R -tree for some profinite ring $R \neq 0$;
- (2) Γ is a \mathbb{Z}_p -tree for some prime p ;
- (3) Γ is a $\mathbb{Z}/p\mathbb{Z}$ -tree for some prime p .

A (connected) profinite graph Γ is called a *pro- p tree* in [41] if it is a $\mathbb{Z}/p\mathbb{Z}$ -tree, or, equivalently, a \mathbb{Z}_p -tree (cf. [43]). The next statement is a slight extension of [6, Theorem 3.9], using precisely the same arguments.

Theorem 9.6. *Let \mathbf{H} be a pseudovariety of groups, F a free pro- \mathbf{H} group on a finite set A and p a prime. Then $\Gamma_A(F)$ is a pro- p tree if and only if every A -generated co-extension of a group in \mathbf{H} by a p -group belongs to \mathbf{H} , or, equivalently, if $G^{\mathbf{Ab}_p} \in \mathbf{H}$ for every A -generated group G in \mathbf{H} .*

Corollary 9.7. *Let \mathbf{H} be a pseudovariety of groups and p a prime. Then the Cayley graph of each finitely generated free pro- \mathbf{H} group is a pro- p tree if and only if $\mathbf{G}_p * \mathbf{H} = \mathbf{H}$.*

The Cayley graphs of the free pro- \mathbf{H} groups of the pseudovariety given in (7.1) are therefore tree-like, but are not pro- p trees for any p .

REFERENCES

1. J. Almeida, *Finite Semigroups and Universal Algebra*, World Scientific, Singapore, 1995. MR **96b**:20069
2. J. Almeida, *Hyperdecidable pseudovarieties and the calculation of semidirect products*, Internat. J. Algebra Comput. **9** (1999), 241–261. MR **2001a**:20102
3. J. Almeida and B. Steinberg, *Iterated semidirect products with applications to complexity*, Proc. London Math. Soc. **80** (2000), 50–74. MR **2000j**:20109
4. J. Almeida and B. Steinberg, *Syntactic and global semigroup theory, a synthesis approach*, pp. 1–23 in: Algorithmic Problems in Groups and Semigroups, J.-C. Birget, S. Margolis, J. Meakin, M. Sapir, eds., Birkhäuser, Basel, 2000. MR **2001c**:20120
5. J. Almeida and P. Weil, *Reduced factorizations in free profinite groups and join decompositions of pseudovarieties*, Internat. J. Algebra Comput. **3** (1994), 375–403. MR **95m**:20066
6. J. Almeida and P. Weil, *Relatively free profinite monoids: An introduction and examples*, pp. 73–117 in: Semigroups, Formal Languages and Groups, J. B. Fountain, ed., Kluwer, Dordrecht, 1995. MR **2000f**:20095
7. J. Almeida and P. Weil, *Free profinite semigroups over semidirect products*, Russian Math. (Iz. VUZ) **39** (1995), 1–27. MR **97e**:20078
8. C. J. Ash, *Inevitable graphs: A proof of the type II conjecture and some related decision procedures*, Internat. J. Algebra Comput. **1** (1991), 127–146. MR **92k**:20114
9. K. Auinger, *Semigroups with central idempotents*, pp. 25–33 in: Algorithmic Problems in Groups and Semigroups, J.-C. Birget, S. Margolis, J. Meakin, M. Sapir, eds., Birkhäuser, Basel, 2000. MR **2001c**:20125
10. K. Auinger, *Join decompositions of pseudovarieties involving semigroups with commuting idempotents*, J. Pure Appl. Algebra **170** (2002), 115–129. MR **2003d**:20082
11. K. Auinger, T. E. Hall, N. R. Reilly and S. Zhang, *Congruences on the lattice of pseudovarieties of finite semigroups*, Internat. J. Algebra Comput. **7** (1997), 433–455. MR **98h**:20100

12. A. H. Clifford and G. B. Preston, *The Algebraic Theory of Semigroups*, Amer. Math. Soc., Providence, 1961. MR **24**:A2627
13. M. Delgado, S. W. Margolis, and B. Steinberg, *Combinatorial group theory, inverse monoids, automata, and global semigroup theory*, Internat. J. Algebra Comput. **12** (2002), 179–211. MR **2003b**:20061
14. S. Eilenberg, *Automata, Languages and Machines*, Academic Press, New York, Vol A, 1974; Vol B, 1976. MR **58**:26604a; MR **58**:26604b
15. G. Elston, *Semigroup expansions using the derived category, kernel, and Malcev products*, J. Pure Appl. Algebra **136** (1999), 231–265. MR **2000a**:20128
16. R. Gitik and E. Rips, *On separability properties of groups*, Internat. J. Algebra Comput. **5** (1995), 703–717. MR **97e**:20059
17. D. Gildenhuys and L. Ribes, *Profinite groups and Boolean graphs*, J. Pure Appl. Algebra **12** (1978), 21–47. MR **81g**:20059
18. M. Hall Jr., *A topology for free groups and related groups*, Ann. of Math. **52** (1950), 127–139. MR **12**:158b
19. B. Herwig and D. Lascar, *Extending partial automorphisms and the profinite topology on free groups*, Trans. Amer. Math. Soc. **352** (2000), 1985–2021. MR **2000j**:20036
20. K. Henckell, S. W. Margolis, J.-E. Pin, and J. Rhodes, *Ash’s type II theorem, profinite topology and Mal’cev products, Part I*, Internat. J. Algebra Comput. **1** (1991), 411–436. MR **93h**:20064
21. P. M. Higgins, *Techniques of Semigroup Theory*, Oxford University Press, Oxford, 1992. MR **93d**:20101
22. J. M. Howie, *Fundamentals of Semigroup Theory*, Clarendon Press, Oxford, 1995. MR **98e**:20059
23. K. Krohn and J. Rhodes, *Algebraic theory of machines, I*, Trans. Amer. Math. Soc. **116** (1965), 450–464. MR **32**:5755
24. K. Krohn and J. Rhodes, *Complexity of finite semigroups*, Ann. of Math. **88** (1968), 128–160. MR **38**:4591
25. M. V. Lawson, *Inverse Semigroups: The Theory of Partial Symmetries*, World Scientific, Singapore, 1998. MR **2000g**:20123
26. S. W. Margolis and J. C. Meakin, *E-unitary inverse monoids and the Cayley graph of a group presentation*, J. Pure Appl. Algebra **58** (1989), 45–76. MR **90f**:20096
27. S. W. Margolis and J. C. Meakin, *Free inverse monoids and graph immersions*, Internat. J. Algebra Comput. **3** (1993), 79–99. MR **94c**:20105
28. S. W. Margolis and J.-E. Pin, *Varieties of finite monoids and topology for the free monoid*, pp 113–130 in: *Proceedings of the 1984 Marquette Conference on Semigroups*, K. Byleen, P. Jones and F. Pastijn, eds., Marquette University, Milwaukee, 1984.
29. S. W. Margolis, M. Sapir, and P. Weil, *Closed subgroups in pro-V topologies and the extension problem for inverse automata*, Internat. J. Algebra Comput. **11** (2001), 405–445. MR **2002g**:20097
30. D. B. McAlister, *Groups, semilattices and inverse semigroups*, Trans. Amer. Math. Soc. **68** (1974), 227–244. MR **50**:10128
31. J. McCammond and J. Rhodes, *Geometric semigroup theory*, Internat. J. Algebra Comput. (to appear).
32. W. D. Munn, *Free inverse semigroups*, Proc. London Math. Soc. **30** (1974), 385–404. MR **50**:13328
33. H. Neumann, *Varieties of Groups*, Springer, Berlin, Heidelberg, New York, 1967. MR **35**:6734
34. G. Niblo, *Separability properties of free groups and surface groups*, J. Pure Appl. Algebra **78** (1992), 77–84. MR **92m**:20019
35. J.-E. Pin, **BG = PG**, *a success story*, pp. 33–47 in: *Semigroups, Formal Languages and Groups*, J. B. Fountain, ed., Kluwer, Dordrecht, 1995. MR **99e**:20072
36. J.-E. Pin and C. Reutenauer, *A conjecture on the Hall topology for the free group*, Bull. London Math. Soc. **23** (1991), 356–362. MR **92g**:20035
37. J. Rhodes and B. Steinberg, *Pointlike sets, hyperdecidability and the identity problem for finite semigroups*, Internat. J. Algebra Comput. **9** (1999), 475–481. MR **2000k**:20075
38. J. Rhodes and B. Steinberg, *Profinite semigroups, varieties, expansions and the structure of relatively free profinite semigroups*, Internat. J. Algebra Comput. **11** (2002), 627–672. MR **2002j**:20114

39. L. Ribes and P. A. Zalesskiĭ, *On the profinite topology on a free group*, Bull. London Math. Soc. **25** (1993), 37–43. MR **93j**:20062
40. L. Ribes and P. A. Zalesskiĭ, *The pro- p topology of a free group and algorithmic problems in semigroups*, Internat. J. Algebra Comput. **4** (1994), 359–374. MR **96e**:20046
41. L. Ribes and P. A. Zalesskiĭ, *Pro- p Trees and Applications*, pp. 75–119 in: New Horizons in pro- p Groups, M. Du Sautoy, D. Segal, A. Shalev, eds., Birkhäuser, Boston, 2000. MR **2001f**:20057
42. L. Ribes and P. A. Zalesskiĭ, *Profinite Groups*, Springer, Berlin, 2000. MR **2001k**:20060
43. L. Ribes and P. A. Zalesskiĭ, *Profinite Trees*, Springer (to appear).
44. J.-P. Serre, *Trees*, Springer-Verlag, Berlin Heidelberg New York, 1980. MR **82c**:20083
45. J. Stallings, *Topology of finite graphs*, Invent. Math. **71** (1983), 551–565. MR **85m**:05037a
46. B. Steinberg, *On pointlike sets and joins of pseudovarieties*, Internat. J. Algebra Comput. **8** (1998), 203–231. MR **99g**:20104
47. B. Steinberg, *On algorithmic problems for joins of pseudovarieties*, Semigroup Forum **62** (2001), 1–40. MR **2002e**:20124
48. B. Steinberg, *Inevitable graphs and profinite topologies: Some solutions to algorithmic problems in monoid and automata theory stemming from group theory*, Internat. J. Algebra Comput. **11** (2001), 25–71. MR **2002a**:68074
49. B. Steinberg, *A delay theorem for pointlikes*, Semigroup Forum **63** (2001), 281–304. MR **2002m**:20094
50. B. Steinberg, *Finite state automata: A geometric approach*, Trans. Amer. Math. Soc. **353** (2001), 3409–3464. MR **2002c**:20106
51. B. Steinberg, *Polynomial closure and topology*, Internat. J. Algebra Comput. **10** (2000), 603–624. MR **2001m**:20102
52. B. Steinberg, *Inverse automata and profinite topologies on a free group*, J. Pure Appl. Algebra **167** (2002), 341–359. MR **2002i**:20037
53. B. Steinberg, **PG** = **BG**: *Redux*, pp. 181–190 in: Proceedings of the International Conference on Semigroups, P. Smith, E. Giraldes, P. Martins, eds., World Scientific, Singapore, 2000. MR **2003b**:20083
54. B. Steinberg, *The uniform word problem for groups and finite Rees quotients of E -unitary inverse semigroups*, J. Algebra **266** (2003), 1–13.
55. J. B. Stephen, *Presentations of inverse monoids*, J. Pure Appl. Algebra **63** (1990), 81–112. MR **91g**:20083
56. B. Tilson, *Categories as algebra: An essential ingredient in the theory of monoids*, J. Pure Appl. Algebra **48** (1987), 83–198. MR **90e**:20061
57. P. Weil, *Computing closures of finitely generated subgroups of the free group*, pp. 289–307 in: Algorithmic Problems in Groups and Semigroups, J.-C. Birget, S. Margolis, J. Meakin, M. Sapir, eds., Birkhäuser, Basel, 2000. MR **2001c**:20049
58. P. A. Zalesskiĭ, *A geometric characterization of free constructions of profinite groups*, Siber. Math. J. **30** (1989), 73–84. MR **90e**:20025
59. P. A. Zalesskiĭ and O. Mel'nikov, *Subgroups of profinite groups acting on trees*, Math. USSR Sbornik **63** (1989), 405–424. MR **90f**:20041

INSTITUT FÜR MATHEMATIK, UNIVERSITÄT WIEN, STRUDLHOFGASSE 4, A-1090 WIEN, AUSTRIA
E-mail address: karl.auinger@univie.ac.at

SCHOOL OF MATHEMATICS AND STATISTICS, CARLETON UNIVERSITY, HERZBERG LABORATORIES,
 OTTAWA, ONTARIO, CANADA K1S 5B6
E-mail address: bsteinbg@math.carleton.ca