

FINITE APPROXIMATION OF FREE GROUPS II: THE THEOREMS OF ASH, HERWIG–LASCAR AND RIBES–ZALESSKII — REVISITED AND STRENGTHENED

K. AUINGER, J. BITTERLICH AND M. OTTO

Dedicated to Prof. Mikhail Volkov on the occasion of his 70th birthday

ABSTRACT. Relations and interactions between the theorems of Ash, Herwig–Lascar and Ribes–Zalesskii are discussed and it is shown that these three theorems are equivalent in the sense that each of them can be derived from each other one. Some strengthening of these theorems that can be obtained by use of the groups provided by the third author’s construction are also considered.

CONTENTS

1. Introduction	2
2. Preliminaries	4
2.1. Inverse monoids	4
2.2. A -generated inverse monoids	5
2.3. Graphs	6
2.4. Cayley graphs of A -generated groups, Schützenberger graphs of A -generated inverse monoids	9
2.5. Extending partial bijections and automorphisms of relational structures	10
2.6. Stallings graphs and Schreier graphs	17
3. Theorems Revisited	19
3.1. Herwig–Lascar implies Ribes–Zalesskii	20
3.2. Ribes–Zalesskii implies Ash	23
3.3. Ash implies Herwig–Lascar	32
3.4. Herwig–Lascar implies Ash	35
3.5. Group-theoretic versions of the Herwig–Lascar Theorem	42
3.6. Construction of groups satisfying Ash’s Theorem	48
4. Theorems Strengthened	52
4.1. Summary of some results from [14, 27, 28]	52
4.2. Proof of Theorem 4.1	60
4.3. Applications	72

2020 *Mathematics Subject Classification.* 05C25, 05E18, 20M18, 20B25, 03C52.

Key words and phrases. permutation group, action graph, Cayley graph, finite inverse monoid, finite n -acyclic group, Theorem of Ash, Theorem of Herwig–Lascar, Theorem of Ribes–Zalesskii.

5. Comparing Theorem 3.3 with Theorem 4.1	76
References	83

1. INTRODUCTION

The three celebrated theorems mentioned in the title were obtained in the 1990s. The Theorem of Ash [6] is about (labellings of graphs by) finite inverse monoids¹, the Theorem of Herwig–Lascar [19] is about relational structures and the Theorem of Ribes–Zalesskii [32] about the profinite topology of the free group. Each of them has had substantial influence in its own field: Ash’s Theorem led to the introduction and investigation of the concepts of *hyperdecidability* and of *tameness* of pseudovarieties of semigroups and monoids [1, 2, 3] initiated by the Porto school and led to many important and interesting questions and results; many authors dealt with extensions and generalisations of the Theorem of Herwig–Lascar (see e.g. [20] and the extensive list of papers cited there); and also the Theorem of Ribes–Zalesskii has been widely discussed: new proofs have been published and the theorem has been generalised to other topologies of the free group and to other groups [33, 10, 12, 38].

From the very beginning various connections and interactions between these theorems have also been discussed: for Ash versus Ribes–Zalesskii this is not surprising since both were originally motivated by giving a solution to the same problem, namely to confirm the so-called *Rhodes type II conjecture*, a conjecture about finite monoids that was much discussed during the 1970s and 1980s. It is known from [37] that the Theorem of Ribes–Zalesskii is equivalent to a special case of Ash’s Theorem (namely to the case of cycle graphs). Likewise, the fact that the Ribes–Zalesskii Theorem is a consequence of the Herwig–Lascar Theorem has been one of the motivating examples in the paper by Herwig and Lascar [19] (the term *profinite topology on free groups* even is displayed in the title). In addition, they presented a version of their theorem — let us call it the *group-theoretic version* of their theorem (Theorem 3.3 in [19]) — which dealt with solvability in the free group of equations of a certain type. This they announced as being equivalent to the *model-theoretic version* of their theorem: Theorem 3.2 in [19], which is usually referred to as *the Herwig–Lascar Theorem*. A special case of the group-theoretic version turns out to be essentially the same as the Ribes–Zalesskii Theorem.

The overall conviction (at least in the semigroup community) is that the Theorem of Ash and the Theorem of Herwig–Lascar are equivalent. Indeed, Almeida and Delgado [4, 5] proved that Ash’s Theorem is equivalent to the above-mentioned group-theoretic version of the theorem: Theorem 3.3 of [19]. But, strictly speaking, a proof that the Theorem of Ash implies

¹By “Ash’s Theorem”, in this paper we always mean its version for inverse monoids.

the model-theoretic version of the Theorem of Herwig–Lascar (Theorem 3.2 of [19]) has — to the best of the authors’ knowledge — not been published, since a proof of the implication

$$\text{group-theoretic version} \implies \text{model-theoretic version}$$

is not published in [19] nor elsewhere. So the equivalence of Ash with Herwig–Lascar is known, but no full proof is published. Ironically, in this context we also have the opposite case: the fact that the Theorem of Ribes–Zalesskii implies the Theorem of Ash (and hence the Theorem of Herwig–Lascar, at least in its group-theoretic version) seems to be not commonly acknowledged or recognised, yet a full proof of this can be retrieved from the literature. In fact, it is one of the key steps in Ash’s paper that the general case of his theorem follows from the special case for cycle graphs (and hence from the Ribes–Zalesskii Theorem).

The first motivation for the present paper is to clean up this kind of mess. In Section 3 we present full proofs for the chain of implications

$$\text{Herwig–Lascar} \implies \text{Ribes–Zalesskii} \implies \text{Ash} \implies \text{Herwig–Lascar}.$$

The most difficult step is the second implication, the first and the third are comparatively easy. The proof of this second step is essentially taken from Ash’s paper, but written up in more detail and somehow polished so that it is accessible, hopefully, also to non-specialists. The method of that proof — which might be called *arborisation of a labelled graph with respect to an inverse monoid* — seems to be interesting in its own right. It will also be used in Section 4. It therefore seems justified to elaborate on the method of that proof in the present paper. The proof of the first implication is inspired by the one in [19]. But, on the one hand, it is enriched with some pointers to the group-theoretic background: it is embedded into the context of Stallings and Schreier graphs, which seem to be *the* natural environment of these arguments. On the other hand, a slightly more general assertion is proved, which prepares the ground for a direct proof of the implication

$$(\text{model-theoretic}) \text{ Herwig–Lascar} \implies \text{Ash}$$

which we also include (Section 3.4). The idea for the third implication is taken from the second author’s doctoral dissertation [17].

The discussion of the subject Ash/Herwig–Lascar/Ribes–Zalesskii (i.e. the content of Section 3) in particular intends to direct the attention of readers interested in model-theoretic extension problems to the relevance of Ash’s work. That part of the paper may also be seen as a kind of survey with the emphasis of shedding new light on the context of these results. In particular, many of the proofs there are in fact new.

Our main motivation, however, to recast the “old stories” lies in the final part of the paper (Section 4). There we present some strengthening of the theory which can be achieved by the application of a construction established by the third author [28, 14]. The main result is Theorem 4.1, a strengthening of Ash’s Theorem. The ideas of the aforementioned (difficult)

proof of the implication Ribes–Zalesskii \implies Ash will be an essential ingredient to establish our result. As an application we also formulate and prove some strengthening of the Theorem of Herwig–Lascar.

In Section 2 we present an extensive aggregation of preliminaries, which should make the part of the paper concerning the old results as self-contained as possible. For the strengthenings obtained in Section 4 we require non-trivial results and concepts from [14]. These are collected in Section 4.1. Finally, in Section 5 some evidence is given that the new Theorem 4.1 is strictly stronger than Ash’s Theorem.

2. PRELIMINARIES

We present a relatively long section of preliminaries to render the paper as self-contained as possible and (hopefully) accessible to non-specialists.

2.1. Inverse monoids. A monoid M is *inverse* if every element $x \in M$ admits a *unique* element x^{-1} , called the *inverse* of x , satisfying $xx^{-1}x = x$ and $x^{-1}xx^{-1} = x^{-1}$. This gives rise to a unary operation $^{-1}: M \rightarrow M$ and an inverse monoid may equivalently be defined as an algebraic structure $(M; \cdot, ^{-1}, 1)$ with \cdot an associative binary operation, 1 a neutral element with respect to \cdot and a unary operation $^{-1}$ satisfying the identities (or laws)

$$(x^{-1})^{-1} = x, (xy)^{-1} = y^{-1}x^{-1}, xx^{-1}x = x \text{ and } xx^{-1}yy^{-1} = yy^{-1}xx^{-1}.$$

In particular, the class of all inverse monoids forms a variety of algebraic structures (in the sense of universal algebra), the variety of all groups $(G; \cdot, ^{-1}, 1)$ being a subvariety. Both, the class of all groups as well as the class of all inverse monoids are varieties of *involutory monoids*, were the variety of all involutory monoids $(M; \cdot, ^{-1}, 1)$ is defined by the identities $(x^{-1})^{-1} = x$ and $(xy)^{-1} = y^{-1}x^{-1}$.

By the Wagner–Preston Theorem [22, Chapter 1, Theorem 1], inverse monoids may as well be characterised as monoids of partial bijections on a set, closed under composition of partial mappings and inversion. Therefore, while groups model symmetries of mathematical structures, inverse monoids (or semigroups) model partial symmetries, that is, symmetries between substructures of mathematical structures. It is therefore not surprising that inverse monoids could play some rôle in the context of extending partial automorphisms of mathematical structures.

Every inverse monoid M is equipped with a partial order \leq , the *natural order*, defined by $x \leq y$ if and only if $x = ye$ for some idempotent e of M (this is equivalent to $x = fy$ for some idempotent f of M). If an inverse monoid M is represented as a monoid of partial bijections, then the idempotents of M are exactly the restrictions of the identity map. Hence any two idempotents commute and therefore the set of all idempotents of M forms a (commutative) submonoid of M . In addition, for every element x of M and idempotent e , the “conjugate” element $x^{-1}ex$ is again an idempotent. For $x, y \in M$ we have $x \leq y$ if and only if $x \subseteq y$, that is, x is a restriction of y .

The order is compatible with the binary operation, that is, from $x \leq y$ and $u \leq v$ it follows that $xu \leq yv$ for all x, y, u, v , and with inversion of M , that is, $x \leq y$ implies $x^{-1} \leq y^{-1}$. For further information on inverse monoids the reader is referred to the monographs by Petrich [25] and Lawson [22].

2.2. A -generated inverse monoids. Throughout, for any non-empty set X (of letters, of edges, etc.) we let $X^{-1} := \{x^{-1} : x \in X\}$ be a disjoint copy of X consisting of formal inverses of the elements of X , and set $\tilde{X} := X \cup X^{-1}$. The mapping $x \mapsto x^{-1}$ is extended to an involution of \tilde{X} by setting $(x^{-1})^{-1} = x$, for all $x \in X$. We let \tilde{X}^* be the free monoid over \tilde{X} , which, subject to the operation $(x_1 \cdots x_n)^{-1} = x_n^{-1} \cdots x_1^{-1}$ (where $x_i \in \tilde{X}$), is the *free involutory monoid* over X . The elements of \tilde{X}^* are called *words over \tilde{X}* , and we identify the empty word with 1 in this context. A word $w \in \tilde{X}^*$ is *reduced* if it does not contain any factor of the form xx^{-1} for $x \in \tilde{X}$. Successive deletion of such factors in a word w until no more such factor is present leads to the *reduced form* $\text{red}(w)$ of w , which is well-defined. A word $w \in \tilde{X}^*$ with $\text{red}(w) = 1$ is called a *Dyck word*.

We fix a non-empty set A (called alphabet in this context). An inverse monoid M together with a (not necessarily injective) mapping $i_M : A \rightarrow M$ (called *assignment function*) is an A -generated inverse monoid if M is generated by $i_M(A)$ as an inverse monoid, that is, generated with respect to the operations $1, \cdot, {}^{-1}$. For every congruence ρ of an A -generated inverse monoid M , the quotient M/ρ is A -generated with respect to the map $i_{M/\rho} = \pi_\rho \circ i_M$ where π_ρ is the projection $M \rightarrow M/\rho$. A *morphism* ψ from the A -generated inverse monoid M to the A -generated inverse monoid N is a homomorphism $M \rightarrow N$ respecting generators from A , that is, satisfying $i_N = \psi \circ i_M$. If it exists, such a morphism is unique and surjective and is called *canonical morphism*, denoted $\psi : M \twoheadrightarrow N$. On a more formal level, an A -generated inverse monoid is an algebraic structure of the form $(M; \cdot, {}^{-1}, 1, A)$ where every symbol $a \in A$ is interpreted in M as a constant (that is, as a nullary operation) via the assignment function i_M . Canonical morphisms of A -generated inverse monoids then are just homomorphisms of algebraic structures in the signature $\{\cdot, {}^{-1}, 1\} \cup A$.

If $M \twoheadrightarrow N$ then M is an *expansion* of N . In our usage, the term “expansion” just concerns the relationship between two individual A -generated inverse monoids M and N . This somehow deviates from the widespread use of that term standing for a functor on certain categories of monoids. The *direct product* of A -generated inverse monoids M_1, \dots, M_k is the A -generated inverse monoid with the diagonal mapping $a \mapsto (i_{M_1}(a), \dots, i_{M_k}(a))$ of A into the full Cartesian product as the natural assignment function; we will denote this product by $M_1 \times \cdots \times M_k$, which in general is a proper submonoid of the full Cartesian product. The special case of A -generated groups will play a significant rôle in this paper. This concerns, in particular, the free

A -generated group which will be denoted F , throughout the paper. The elements of F will usually be identified with the reduced words over \tilde{A}^* .

As already mentioned, the assignment function is not necessarily injective, and, what is more, some generators may even be sent to the identity element of M . This is not a deficiency, but rather is adequate in our context, since we want the quotient of an A -generated structure to be again A -generated.

The assignment function i_M is usually not explicitly mentioned; it uniquely extends to a homomorphism $[\]_M: \tilde{A}^* \rightarrow M$ (of involutory monoids). For every word $p \in \tilde{A}^*$, $[p]_M$ is the *value of p in M* or simply the *M -value of p* . For two words $p, q \in \tilde{A}^*$, the A -generated inverse monoid M satisfies the relation $p = q$ if $[p]_M = [q]_M$, in which case the words p and q are *M -related* or *M -equivalent*, while M avoids the relation $p = q$ if $[p]_M \neq [q]_M$.

2.3. Graphs. In this paper, we consider the Serre definition [34] of graph structures, admitting multiple directed edges between pairs of vertices and directed loop edges at individual vertices. In the literature, such structures are often called *multidigraphs*, *directed multigraphs* or *quivers*. The following formalisation is convenient for our purposes. A *graph* \mathcal{K} is a structure of the form $(V \cup K; \alpha, \omega, {}^{-1})$ with V its set of *vertices*, K its set of *edges* (disjoint from V), with *incidence functions* $\alpha: K \rightarrow V$ and $\omega: K \rightarrow V$, specifying, for each edge e (including loop edges), the *initial* vertex αe and the *terminal* vertex ωe , and *involution* ${}^{-1}: K \rightarrow K$ satisfying $\alpha e = \omega e^{-1}$, $\omega e = \alpha e^{-1}$ and $e \neq e^{-1}$ for every edge $e \in K$. Instead of *initial/terminal vertex* the terms *source/target* are also used in the literature. One should think of an edge e with initial vertex $\alpha e = u$ and terminal vertex $\omega e = v$ in “geometric” terms as $e: \underset{u}{\bullet} \longrightarrow \underset{v}{\bullet}$ and of its inverse $e^{-1}: \underset{u}{\bullet} \longleftarrow \underset{v}{\bullet}$ as “the same edge but traversed in the opposite direction”. The pair $\{e, e^{-1}\}$ consisting of an edge e and its inverse e^{-1} is sometimes called a *geometric edge* of the graph \mathcal{K} and is denoted e^\bullet , the set of all geometric edges is denoted K^\bullet . The *degree* of a vertex v in a graph \mathcal{K} is the number of geometric edges in \mathcal{K} incident with v ; this is the same as the number of edges e having initial vertex $\alpha e = v$ and also the number of edges e having terminal vertex $\omega e = v$. For a given graph \mathcal{K} we sometimes denote the set of its vertices by $V(\mathcal{K})$, the set of its edges by $K(\mathcal{K})$. A graph $(V \cup K; \alpha, \omega, {}^{-1})$ is *oriented* if the edge set K is partitioned as $K = E \cup E^{-1} = \tilde{E}$ such that every ${}^{-1}$ -orbit contains exactly one element of E and one of E^{-1} ; the edges in E are the *positive* or *positively oriented* edges, those in E^{-1} the *negative* or *negatively oriented* ones.

A *subgraph* of the graph \mathcal{K} is a substructure that is induced over a subset of $V \cup K$ which is closed under the operations α and ${}^{-1}$ (and therefore also under ω). In particular, every non-empty subset $S \subseteq V \cup K$ generates a unique minimal subgraph $\langle S \rangle$ of \mathcal{K} containing S , which is the *subgraph of \mathcal{K} spanned by S* . Likewise, for graphs $\mathcal{K} = (V \cup K; \alpha, \omega, {}^{-1})$ and $\mathcal{L} = (W \cup L; \alpha, \omega, {}^{-1})$ a *morphism* $\varphi: \mathcal{K} \rightarrow \mathcal{L}$ is a mapping $\varphi: V \cup K \rightarrow W \cup L$ mapping vertices to vertices, edges to edges that is compatible with the operations α and ${}^{-1}$

(and therefore also ω). A *congruence* Θ on the graph $\mathcal{K} = (V \cup K; \alpha, \omega, {}^{-1})$ is an equivalence relation $\Theta = \Theta_V \cup \Theta_K$ with Θ_V and Θ_K being equivalence relations on V and K , respectively, compatible with the operations α and ${}^{-1}$ (therefore also ω):

$$e \Theta_K f \implies \alpha e \Theta_V \alpha f, \omega e \Theta_V \omega f, e^{-1} \Theta_K f^{-1} \text{ for all } e, f \in K$$

and the concept of the quotient graph \mathcal{K}/Θ is the obvious one, with the usual homomorphism theorem. A *non-empty path* π in \mathcal{K} is a sequence $\pi = e_1 e_2 \cdots e_n$ ($n \geq 1$) of *consecutive* edges (that is $\omega e_i = \alpha e_{i+1}$ for all $1 \leq i < n$); we set $\alpha\pi := \alpha e_1$ and $\omega\pi := \omega e_n$ (denoting the initial and terminal vertices of the path π); the *inverse path* π^{-1} is the path $\pi^{-1} := e_n^{-1} \cdots e_1^{-1}$; it has initial vertex $\alpha\pi^{-1} = \omega\pi$ and terminal vertex $\omega\pi^{-1} = \alpha\pi$. A path π is *closed* or a *cycle* [at the vertex v] if $\alpha\pi = \omega\pi [= v]$. A path is *simple* if no vertex occurs twice, except that possibly the initial and terminal vertices may coincide, in which case we have a simple closed path. We also consider, for each vertex v , the *empty path at v* , denoted ε_v for which we set $\alpha\varepsilon_v = v = \omega\varepsilon_v$ and $\varepsilon_v^{-1} = \varepsilon_v$ (it is convenient to identify ε_v with the vertex v itself). We say that π is a path *from* $u = \alpha\pi$ *to* $v = \omega\pi$, and we will also say that u and v are *connected by* π (and likewise by π^{-1}). A path π is *reduced* if it does not contain any segment of the form ee^{-1} ; successive deletion of such segments until there is no more such segment leads to the *reduced form* $\text{red}(\pi)$ of π , which is again a path and has the same initial and terminal vertices as the original path π .

A graph is *connected* if any two vertices can be connected by some path. A connected graph \mathcal{T} is a *tree* if it does not admit non-empty reduced closed paths; a graph is a *forest* if each connected component is a tree. The subgraph $\langle \pi \rangle$ spanned by the non-empty path π is the graph spanned by the edges of π ; it coincides with $\langle \pi^{-1} \rangle$; the graph spanned by an empty path ε_v simply is $\{v\}$ (one vertex, no edge). An *arc* is a graph spanned by a simple non-closed path; a *cycle* or a *cycle graph* is a graph spanned by a simple closed path. The *length* of a cycle or an arc is the number of its geometric edges. The *cyclomatic number* of a finite graph $\mathcal{K} = (K \cup V; \alpha, \omega, {}^{-1})$ is $|K^\bullet| - |V| + c$ where c is the number of the connected components of \mathcal{K} . This number is non-negative and is equal to 0 if and only if \mathcal{K} is a forest.

Let A be a finite set; a *word labelling over the alphabet A* of the graph $\mathcal{K} = (V \cup K; \alpha, \omega, {}^{-1})$ is a mapping $\ell: K \rightarrow \tilde{A}^*$ respecting the involution: $\ell(e^{-1}) = \ell(e)^{-1}$ for all $e \in K$. If every label $\ell(e)$ is an element of \tilde{A} we call it an *A-labelling*. Every *A-labelling* $\ell: K \rightarrow \tilde{A}$ gives rise to an *orientation* of \mathcal{K} : setting $E := \{e \in K: \ell(e) \in A\}$ (positive edges) and $E^{-1} := \{e \in K: \ell(e) \in A^{-1}\}$ (negative edges), it follows that $E \cap E^{-1} = \emptyset$ and we get $K = \tilde{E}$. For a path $e_1 \cdots e_k$ in a word-labelled graph over \tilde{A} as well as in an *A-labelled* graph \mathcal{K} , its label is $\ell(e_1 \cdots e_k) := \ell(e_1) \cdots \ell(e_k)$, which is a word over \tilde{A} .

We consider A -labelled graphs as structures $\mathcal{K} = (V \cup K; \alpha, \omega, {}^{-1}, \ell, A)$ in their own right. By a *subgraph of an A -labelled graph* we mean just a subgraph with the induced labelling. Morphisms $\varphi: \mathcal{K} \rightarrow \mathcal{L}$ of A -labelled graphs and congruences Θ on an A -labelled graph \mathcal{K} are defined accordingly. They must be morphisms, respectively congruences, on the underlying graphs and must in addition respect the labelling. For a morphism $\varphi: \mathcal{K} \rightarrow \mathcal{L}$ this means that $\ell(\varphi(e)) = \ell(e)$ for every edge e of \mathcal{K} . For a congruence Θ on \mathcal{K} this means that $e \Theta f \implies \ell(e) = \ell(f)$ for all edges $e, f \in \mathcal{K}$. This guarantees that the labelling ℓ is well-defined on the quotient graph \mathcal{K}/Θ . Again we have the usual homomorphism theorem.

An A -labelled graph is an A -graph if every vertex u has, for every label $a \in \tilde{A}$, at most one edge with initial vertex u and label a . In the literature, such graphs occur under a variety of different names, such as *folded graph* [21] or *inverse automaton* [9, 12], to mention just two. Every A -labelled graph \mathcal{K} admits a unique smallest congruence Θ such that the quotient \mathcal{K}/Θ is an A -graph. In case \mathcal{K} is finite, this quotient can be computed nicely by the following procedure: if there is a vertex from which there emerge two edges e_1, e_2 having the same label $a \in \tilde{A}$, then identify these edges and their

terminal vertices: $\begin{array}{c} a \nearrow \bullet \\ a \searrow \bullet \end{array} \rightsquigarrow \bullet \xrightarrow{a} \bullet$, and continue this process until no such

pair of edges is left. This algorithm is called *Stallings foldings* and goes back to Stallings [35]. It is useful to have in mind that two vertices p_1 and p_2 are identified by this process (that is, by the congruence Θ) if and only if there is a path $p_1 \longrightarrow p_2$ labelled by a Dyck word. It follows that in an A -graph, if the label $\ell(\pi)$ of a path π is a Dyck word over A then the path itself is a Dyck word over E (with E the set of positive edges of \mathcal{K} , the orientation of \mathcal{K} being induced by the A -labelling). In an A -graph \mathcal{K} , for every word $p \in \tilde{A}^*$ and every vertex u there is at most one path $\pi = \pi_u^{\mathcal{K}}(p)$ with initial vertex $\alpha\pi = u$ and label $\ell(\pi) = p$. We set $u \cdot p := \omega(\pi_u^{\mathcal{K}}(p))$, the terminal vertex of the path starting at u and being labelled p if such a path p exists; then, for every $p \in \tilde{A}^*$, the partial mapping $[p]: V \rightarrow V$, $u \mapsto u \cdot p$ is a partial bijection of the vertex set V of \mathcal{K} . Thus the involutory monoid \tilde{A}^* acts on V by partial bijections on the right. The inverse monoid

$$\mathcal{T}(\mathcal{K}) := \{[p]: p \in \tilde{A}^*\} \quad (2.1)$$

of partial bijections obtained this way, is called the *transition inverse monoid* $\mathcal{T}(\mathcal{K})$ of the A -graph \mathcal{K} . We note that the empty word $1 \in \tilde{A}^*$ induces the identity mapping on all of V . The transition inverse monoid $\mathcal{T}(\mathcal{K})$ is an A -generated inverse monoid in a natural way, the letter $a \in \tilde{A}$ induces the partial bijection $[a]$ on the vertex set that maps every vertex u to the terminal vertex $\omega\pi_u(a)$ of the edge $\pi_u(a)$, which is the unique edge with initial vertex u and label a if such an edge exists.

Similarly as for A -generated inverse monoids, morphisms of A -graphs are essentially unique: given two connected A -graphs \mathcal{K} and \mathcal{L} , there is, for any

pair of vertices $u \in V(\mathcal{K})$ and $v \in V(\mathcal{L})$, at most one morphism $\mathcal{K} \rightarrow \mathcal{L}$ of A -graphs mapping u to v . In case such a morphism exists and is surjective we address it as *canonical morphism* $\mathcal{K} \twoheadrightarrow \mathcal{L}$.

An A -graph \mathcal{K} is called *complete* or a *group action graph* (also called *permutation automaton*) if every vertex u has, for every label $a \in \tilde{A}$ exactly one edge f with initial vertex $\alpha f = u$ and label $\ell(f) = a$. In this case, for every word $p \in \tilde{A}^*$ and every vertex u there exists exactly one path $\pi = \pi_u^{\mathcal{K}}(p)$ starting at u and having label p . Hence the transition inverse monoid $\mathcal{T}(\mathcal{K})$ is a group of permutations of V .

2.4. Cayley graphs of A -generated groups, Schützenberger graphs of A -generated inverse monoids. Given an A -generated group Q , the *Cayley graph* \mathcal{Q} of Q has vertex set Q and, for $g, h \in Q$ and $a \in \tilde{A}$ there is an edge $g \rightarrow h$ labelled a if and only if $g[a]_Q = h$. More formally we define the Cayley graph \mathcal{Q} of the A -generated group Q as follows:

- the set of vertices of \mathcal{Q} is Q ,
- the set of edges of \mathcal{Q} is $Q \times \tilde{A}$, and, for $g \in Q$, $a \in \tilde{A}$, the incidence functions, involution and labelling are defined according to

$$\begin{aligned} \alpha(g, a) &:= g, \\ \omega(g, a) &:= g[a]_Q, \\ (g, a)^{-1} &:= (g[a]_Q, a^{-1}), \\ \ell(g, a) &:= a. \end{aligned}$$

The edge (g, a) should be thought of as $\bullet \xrightarrow[g]{a} \bullet$, its inverse as $\bullet \xleftarrow[ga]{a^{-1}} \bullet$, where ga stands for $g[a]_Q$. We note that Q acts on \mathcal{Q} by left multiplication as a group of automorphisms via

$$g \mapsto {}^h g := hg \quad \text{and} \quad (g, a) \mapsto {}^h(g, a) := (hg, a)$$

for all $g, h \in Q$ and $(g, a) \in Q \times \tilde{A}$, where h , as a group element, acts on the vertex g of \mathcal{Q} and on the edge (g, a) of \mathcal{Q} . Clearly, the Cayley graph \mathcal{Q} is not an invariant of the plain group Q but depends on the assignment function i_Q . The group Q itself is isomorphic to the automorphism group of its Cayley graph viewed as an A -graph.

One could equally well define the Cayley graph of an A -generated inverse monoid M . It has turned out however, that a collection of certain subgraphs of that graph — the Schützenberger graphs — is more important than the entire Cayley graph itself. Recall that Green's equivalence relation \mathcal{R} is defined on any monoid M by setting, for $m, n \in M$, $m \mathcal{R} n$ if and only if $mM = nM$ (that is, m and n generate the same right ideal of M). The \mathcal{R} -class of an element m is usually denoted R_m . For any inverse monoid M , we have that $m \mathcal{R} n$ if and only if $mm^{-1} = nn^{-1}$. Equivalently, $m \mathcal{R} n$ if and only if $m = nx$ and $n = mx^{-1}$ for some x . In particular, every \mathcal{R} -class R_m of M contains exactly one idempotent, namely mm^{-1} . Schützenberger graphs

have been introduced by Stephen [36] and are defined as follows. For an A -generated inverse monoid M and an element $m \in M$ the *Schützenberger graph* \mathcal{M}_m of M with respect to m has vertex set R_m and for $p, r \in R_m$ and $a \in \tilde{A}$ there is an edge $p \longrightarrow r$ labelled a if and only if $p[a]_M = r$. Important for us is the following result.

Lemma 2.1 (Theorem 3.2 in [36]). *For every A -generated inverse monoid M and any $m \in M$, the Schützenberger graph \mathcal{M}_m is an A -graph, and for any word $w \in \tilde{A}^*$ we have $[w]_M \geq m$ if and only if w labels a path $mm^{-1} \longrightarrow m$ in \mathcal{M}_m .*

2.5. Extending partial bijections and automorphisms of relational structures. We are concerned with a (usually) finite set X and a set of partial bijections on X , usually denoted A , and we are interested in a superset $Y \supseteq X$ and, for every $a \in A$ a bijection \hat{a} of Y that extends a . Later we shall assume that X and Y carry relational structure and the involved (partial) mappings are (partial) automorphisms of the structure in question. The mappings $a \in A$ and their extensions \hat{a} (often also simply denoted a) are assumed to act on the right, written as $x \mapsto x \cdot a$ (partial mappings) or $x \mapsto xa$ (total mappings).

In the paper it will be necessary to consider certain algebraic structures generated by a set of symbols whose elements are in bijective correspondence with the elements of A ; likewise we will need to consider graphs whose edges are labelled by such symbols. In order to avoid overly cumbersome notation we will not distinguish notationally between such symbols and partial bijections on X (just as in [19]). Hence A and $a \in A$ will denote both, depending on the context: a set of partial bijections on a set X and at the same time a set of symbols. In both cases we denote by \tilde{A} the set $A \cup A^{-1}$, and we note that in case of abstract symbols we always have $A \cap A^{-1} = \emptyset$, while in case of partial bijections, an element a may be self inverse: $a \in A \cap A^{-1}$.

2.5.1. Embedding A -graphs into complete A -graphs and unary algebras. Suppose we are given a finite set X and a set A of partial bijections on X . The data (X, \tilde{A}) can be encoded in an A -graph \mathcal{P} in an obvious way: the vertices are the elements of X and, for every pair $x_1, x_2 \in X$, there is an edge $e: x_1 \longrightarrow x_2$ with label $\ell(e) = a \in \tilde{A}$ if and only if $x_1 \cdot a = x_2$. Note that $x_1 \cdot a = x_2$ if and only if $x_2 \cdot a^{-1} = x_1$. Connected with this is the transition inverse monoid $\mathcal{T}(\mathcal{P})$ whose action will be denoted $x \mapsto x \cdot [w]_{\mathcal{T}(\mathcal{P})}$ for the partial bijection $[w]_{\mathcal{T}(\mathcal{P})}$ and $w \in \tilde{A}^*$. To have a superset $Y \supseteq X$ and bijections \hat{a} on Y that extend the partial bijections $a \in \tilde{A}$ is the same as to have an embedding $\mathcal{P} \hookrightarrow \bar{\mathcal{P}}$ of the A -graph \mathcal{P} into a complete A -graph $\bar{\mathcal{P}}$ on the vertex set Y . In this situation, the complete graph $\bar{\mathcal{P}}$ is *some completion* of \mathcal{P} . While there are many ways to complete an incomplete A -graph \mathcal{P} , there is an obviously unique maximal one, which is infinite. In order to construct it, take the Cayley graph \mathcal{F} of the free A -generated group F , choose some letter $a \in \tilde{A}$ and remove from the graph all edges e (and their inverses) with

initial vertex 1_F whose label is not a . The graph \mathcal{F} splits into $2|A|$ subtrees of \mathcal{F} , exactly one of which contains the vertex 1_F . Call this tree the *free a -cone* $[a\mathcal{F}]$ with base vertex 1_F , now re-denoted 1_a . Next take the incomplete A -graph \mathcal{P} and attach a free a -cone $[a\mathcal{F}]$ to every vertex $v \in V(\mathcal{P})$ that has no edge e with this initial vertex $\alpha e = v$ and label $\ell(e) = a$. Formally, this is achieved by forming the disjoint union $\mathcal{P} \sqcup [a\mathcal{F}]$ and factoring by the congruence identifying the vertices v and 1_a and leaving all other vertices and edges unidentified. By doing so for all vertices $v \in V(\mathcal{P})$ and all letters $a \in \tilde{A}$ we get a complete A -graph. We denote this graph as \mathcal{PF} and call it the *free completion of \mathcal{P}* .

A more intuitive approach to get the free completion \mathcal{PF} of an A -graph \mathcal{P} is as follows. For every vertex v and every letter $a \in \tilde{A}$ for which there is no edge in \mathcal{P} with initial vertex v and label a , add a new vertex va , say, and a new edge $e: v \xrightarrow{a} va$, and the inverse of this new edge. Every newly added vertex requires the addition of $2|A| - 1$ further vertices; indeed, only one edge emerges from va , namely $e^{-1}: va \xrightarrow{a^{-1}} v$ carrying the label a^{-1} . Hence, for every $b \in \tilde{A} \setminus \{a^{-1}\}$ we have to further add a new vertex vab , say, and the edge $va \xrightarrow{b} vab$ (and its inverse), and we have to continue to do so ad infinitum. Having done so, we have attached a free a -cone $[a\mathcal{F}]$ to the vertex v . The result \mathcal{PF} may thus be obtained as the union $\bigcup \mathcal{P}_n$ of an ascending chain $\mathcal{P} \subseteq \mathcal{P}_1 \subseteq \mathcal{P}_2 \subseteq \dots \subseteq \mathcal{P}_n \subseteq \dots$ of extensions of \mathcal{P} where $\mathcal{P}_n \supseteq \mathcal{P}$ is comprised of the vertices of \mathcal{PF} which have distance at most n from \mathcal{P} and the edges (and their inverses) required to access them from \mathcal{P} .

The process of successively adding vertices va_1, va_1a_2, \dots leads to the universal algebraic view of the situation. The letters $a \in \tilde{A}$ may be viewed as *algebraic operations* successively applied to the elements v of the given *generating set* X (the vertices of \mathcal{P}). To make this more precise, let $Y_{\mathcal{P}} := V(\mathcal{PF})$ be the set of vertices of \mathcal{PF} ; every letter $a \in \tilde{A}$ induces a bijective mapping $y \mapsto ya$ on $Y_{\mathcal{P}}$. Hence $(Y_{\mathcal{P}}; \tilde{A})$ may be interpreted as a *unary algebra* that is, an algebraic structure all of whose fundamental operations are unary functions. Let us call any unary algebra $(Z; \tilde{A})$ whose fundamental operations are denoted by the elements of \tilde{A} an *\tilde{A} -set*. In the context of \tilde{A} -sets $(Z; \tilde{A})$ we shall use the following notation: for a word $w = a_1 \dots a_n \in \tilde{A}^*$ and $z \in Z$ set

$$zw := (((\dots((za_1)a_2)\dots)a_{n-1})a_n).$$

Words $w \in \tilde{A}$ of length at least 2 do not belong to the algebraic signature but can be seen as *term operations* on \tilde{A} -sets. We mention two important features of the \tilde{A} -set $(Y_{\mathcal{P}}; \tilde{A})$ obtained from \mathcal{PF} in this way:

- (i) it is generated, as an \tilde{A} -set, by the subset $X = V(\mathcal{P})$ of $Y_{\mathcal{P}}$ (for every $y \in Y_{\mathcal{P}}$ there exists $x \in X$ and $w \in \tilde{A}^*$ such that $y = xw$, that is, $Y_{\mathcal{P}}$ is the closure of X under the operations \tilde{A} of $(Y_{\mathcal{P}}; \tilde{A})$), and

- (ii) it satisfies every identity (or law) of the form $yw = y$ for every Dyck word w .

Let us call \tilde{A} -sets which satisfy (ii) *F-compatible \tilde{A} -sets*. For later use we note that every *F-compatible \tilde{A} -set* $(Y; \tilde{A})$ gives rise to a complete *A-graph* whose set of vertices is Y , and for $y_1, y_2 \in Y$ and $a \in \tilde{A}$ there is an edge $y_1 \longrightarrow y_2$ having label a if and only if $y_1 a = y_2$. We call this graph the *A-graph associated with the F-compatible \tilde{A} -set* $(Y; \tilde{A})$ or the *functional graph* of $(Y; \tilde{A})$.

The class of all *F-compatible \tilde{A} -sets* forms a variety in the sense of universal algebra, hence this class admits free objects. A realisation (or model) of the free *F-compatible \tilde{A} -set* generated by X is given by the base set $X \times F$ endowed with the operations $(x, g)a := (x, g[a]_F)$ for every $x \in X$, $g \in F$ and $a \in \tilde{A}$, and X is embedded in $X \times F$ via $x \mapsto (x, 1_F)$. It will be convenient to denote the base set $X \times F$ of this free *F-compatible \tilde{A} -set* simply by XF ; for its elements we write xg for (x, g) (if $g \neq 1_F$) and just x for $(x, 1_F)$. We see that the X -generated *F-compatible \tilde{A} -set* $(Y_{\mathcal{P}}; \tilde{A})$, which arises from the free completion $\mathcal{P}F$ of the graph \mathcal{P} , is not a free *F-compatible \tilde{A} -set* (unless every $a \in \tilde{A}$ represents the empty transformation, or, in other words, the graph \mathcal{P} which encodes the data (X, \tilde{A}) has no edges). Conversely, this means that the edges of the graph \mathcal{P} , or, equivalently, the partial mappings $a \in \tilde{A}$, if non-empty, can be interpreted as *non-trivial relations* which provide a *presentation* of the X -generated *F-compatible \tilde{A} -set* $(Y_{\mathcal{P}}; \tilde{A})$:

$$(Y_{\mathcal{P}}; \tilde{A}) \cong (XF; \tilde{A}) / \Theta(\mathcal{P})$$

where $\Theta(\mathcal{P})$ is the smallest \tilde{A} -set congruence on $(XF; \tilde{A})$ containing all pairs $(x_1[a]_F, x_2)$ for $x_1, x_2 \in X$ and $a \in \tilde{A}$ for which $x_1 \xrightarrow{a} x_2$ is an edge of \mathcal{P} . Thus the *F-compatible \tilde{A} -set* $(XF; \tilde{A}) / \Theta(\mathcal{P})$ is the *F-compatible \tilde{A} -set defined by the generating set X subject to the defining relations \mathcal{P}* and will be denoted $(X_{\mathcal{P}}F; \tilde{A})$. The *A-graph* associated with this freely presented *F-compatible \tilde{A} -set* $(X_{\mathcal{P}}F; \tilde{A})$ is then just the free extension $\mathcal{P}F$ we started with.

In the following, we will no longer distinguish notationally between the base set of an algebraic structure and the algebraic structure itself. This is usual when dealing with algebraic structures and we did do so for groups, inverse monoids, etc.

We may generalise the concept of an *F-compatible \tilde{A} -set*. Instead of F take any *A-generated group* G ; an *F-compatible \tilde{A} -set* $(Y; \tilde{A})$ (or just Y) is called a *G-compatible \tilde{A} -set* if it satisfies all identities of the form $yw = y$ for all $w \in \tilde{A}^*$ for which $[w]_G = 1_G$. Again the class of all *G-compatible \tilde{A} -sets* forms a variety of \tilde{A} -sets and hence admits free objects. Again the free *G-compatible \tilde{A} -set* generated by X , denoted XG , can be realised as the base set $X \times G$ endowed with the operations $(x, g)a := (x, g[a]_G)$ for all

$x \in X$, $g \in G$ and $a \in \tilde{A}$. The elements of XG will again be denoted as xg etc. (instead of (x, g) etc.). And again the relations given by the partial mappings \tilde{A} on X or by its graph realisation \mathcal{P} , respectively, provide non-trivial relations which give rise to the G -compatible \tilde{A} -set $X_{\mathcal{P}}G$ generated by X subject to the defining relations \mathcal{P} which is defined as

$$X_{\mathcal{P}}G := XG/\Theta(\mathcal{P}) \quad (2.2)$$

where now $\Theta(\mathcal{P})$ is the smallest \tilde{A} -set congruence on XG containing all pairs $(x_1[a]_G, x_2)$ for all $x_1, x_2 \in X$ and $a \in \tilde{A}$ for which $x_1 \xrightarrow{a} x_2$ is an edge in \mathcal{P} .

In this construction it may happen that distinct elements $x_1, x_2 \in X$ are related by $\Theta(\mathcal{P})$. Obviously, two distinct elements $x_1, x_2 \in X$ will be related by $\Theta(\mathcal{P})$ if there is a path $\pi: x_1 \rightarrow x_2$ in \mathcal{P} whose label $\ell(\pi)$ has G -value 1_G . That it happens only in this case follows from the next lemma which may be seen as a solution to the word problem of the presentation of the algebraic structure $X_{\mathcal{P}}G$.

Lemma 2.2. *Let \mathcal{P} be an A -graph with vertex set X and let G be an A -generated group. Then $xg \Theta(\mathcal{P}) yh$ holds for $xg, yh \in XG$ if and only if there exists a word $w \in \tilde{A}^*$ such that $x \cdot [w]_{\mathcal{T}(\mathcal{P})} = y$ and $gh^{-1} = [w]_G$.*

Proof. Consider the binary relation Υ on XG defined by

$$xg \Upsilon yh : \Longleftrightarrow \exists w \in \tilde{A}^* : x \cdot [w]_{\mathcal{T}(\mathcal{P})} = y \text{ and } gh^{-1} = [w]_G.$$

It is routine to verify that (i) $\Upsilon \subseteq \Theta(\mathcal{P})$, (ii) Υ is an equivalence relation and (iii) Υ is compatible with the \tilde{A} -set structure, that is, it is a congruence. Since $\Theta(\mathcal{P})$ is the smallest congruence which contains all pairs $(x[w]_G, y)$ for which the equality $x \cdot [w]_{\mathcal{T}(\mathcal{P})} = y$ holds, it follows that $\Upsilon = \Theta(\mathcal{P})$. \square

Of particular importance is when no two distinct elements $x_1, x_2 \in X$ are related under $\Theta(\mathcal{P})$; that is, $\Theta(\mathcal{P}) \upharpoonright X = \text{id}_X$. For later use we state this concept as a definition.

Definition 2.3. An A -generated group G extends an A -graph \mathcal{P} if

$$\Theta(\mathcal{P}) \upharpoonright V(\mathcal{P}) = \text{id}_{V(\mathcal{P})}.$$

The following, related concept will also be used throughout the paper.

Definition 2.4. An A -generated group G extends an A -generated inverse monoid M if, for all words $w \in \tilde{A}$, $[w]_G = 1_G$ implies that $[w]_M$ is idempotent.

In semigroup-theoretic terms, an A -generated group G extends an A -generated inverse monoid M in the sense of Definition 2.4 if and only if the A -generated product $M \times G$ is an E -unitary cover of M . It is clear that G extends the A -graph \mathcal{P} in the sense of Definition 2.3 if and only if G extends the transition inverse monoid $\mathcal{T}(\mathcal{P})$ in the sense of Definition 2.4. The value $[w]_M$ of any Dyck word w in an A -generated inverse monoid M

is an idempotent — this is easily seen by induction on the length w and taking into account that $x^{-1}ex$ is an idempotent for every element x and every idempotent e . Therefore, the free group F extends every A -generated inverse monoid M . The existence of a *finite* A -generated group G extending a given finite A -generated inverse monoid M is less obvious, but still well-known and easy to see: take any representation of M as a monoid of partial bijections on some finite set X , take any extensions to total bijections on X (or on some finite superset Y of X) of the partial bijections represented by the elements of A and consider the group G of permutations generated by these bijections — then any such group G of permutations of X (or of Y) extends M in the sense of Definition 2.4. In particular, for a given finite A -generated inverse monoid M there is no canonical choice for a finite A -generated group G to extend M — but there exist *definable* choices for given M . By this we mean that we could single out one of several concrete specifications that yield an unambiguous, explicitly definable output $G(M, A; \iota)$ for any given M , depending just on M and the assignment of generators $\iota: a \mapsto [a]_M$ in an isomorphism-invariant manner even up to permutations of the set A (in the sense that $(M, A; \iota) \cong (M', A; \iota')$ implies $G(M, A; \iota) \cong G(M', A; \iota')$). One obvious such specification uses the above recipe based on the unique natural representation of M as a monoid of partial bijections on the set M itself, and the direct product of all the extensions to groups that arise from different combinations of extensions of the partial to global bijections of $X := M$. (This is but one manner of replacing all arbitrary choices in the general recipe by explicitly definable choices and/or parallelisation of all possible choices.) The potentially interesting advantage of such a concrete and explicit specification with definable outcome precisely lies in its isomorphism-invariant nature, which in particular guarantees compatibility with internal symmetries. If, for instance, $[a]_M$ and $[a']_M$ are related by an automorphism of $(M, A; \iota)$, then that translates, just by isomorphism-invariance, into the existence of an automorphism of such G that correspondingly relates $[a]_G$ and $[a']_G$. This fundamental feature, which is treated in some detail in [29], relies on explicit definability, while ‘free choices’ as in the general recipe above can break internal symmetries (so that, e.g., $[a]_G$ and $[a']_G$ could have different order in G .)

Also note that G extends \mathcal{P} in the sense of Definition 2.3 if and only if the natural mapping $X \rightarrow X_{\mathcal{P}}G$, $x \mapsto x\Theta(\mathcal{P})$ provides an embedding of the A -graph \mathcal{P} into the complete A -graph associated with the \tilde{A} -set $X_{\mathcal{P}}G$ mentioned earlier. An immediate consequence of Lemma 2.2 is the following well-known fact.

Lemma 2.5. *Let \mathcal{P} be an \tilde{A} -graph on vertex set X and let G be an A -generated group; then G extends \mathcal{P} if and only if for each word $w \in \tilde{A}^*$ for which $[w]_G = 1_G$ the partial transformation $[w]_{\mathcal{T}(\mathcal{P})}$ on X is a partial identity, that is, $x \cdot [w]_{\mathcal{T}(\mathcal{P})} = x$ for all $x \in X$ for which it is defined. This is*

equivalent to the fact that $[w]_{\mathcal{T}(\mathcal{P})}$ is an idempotent in the transition inverse monoid $\mathcal{T}(\mathcal{P})$.

2.5.2. Relational structures. We thus are ready for the next step. We now think of X as the universe of a relational structure, i.e. as being equipped with interpretations of the relation symbols in a fixed finite relational signature σ , and of the given partial bijections \tilde{A} of X as *partial automorphisms* of that structure. So, let us fix a finite relational signature σ consisting of symbols R , say, such that each such R has an *interpretation* as a relation $R^{\mathfrak{X}}$ on X , of some arity r . That is, every symbol R comes with a well-defined subset $R^{\mathfrak{X}} \subseteq X^r$ (r the arity of R); the structure $(X; (R^{\mathfrak{X}})_{R \in \sigma})$ will be denoted \mathfrak{X} . For $Z \subseteq X$, a σ -structure $\mathfrak{Z} = (Z; (R^{\mathfrak{Z}})_{R \in \sigma})$ is a *weak substructure* of \mathfrak{X} if $R^{\mathfrak{Z}} \subseteq R^{\mathfrak{X}}$ for all $R \in \sigma$, while \mathfrak{Z} is an *induced substructure* of \mathfrak{X} if $R^{\mathfrak{Z}} = R^{\mathfrak{X}} \cap Z^r$ for all $R \in \sigma$ where r denotes the arity of R . A σ -structure \mathfrak{Y} is an *extension* of \mathfrak{X} if \mathfrak{X} is an induced substructure of \mathfrak{Y} .

We are interested in extensions \mathfrak{Y} of \mathfrak{X} such that the partial automorphisms of \mathfrak{X} extend to total automorphisms of \mathfrak{Y} . In this context it is convenient to introduce some further terminology. Suppose we are given a σ -structure $\mathfrak{X} = (X; (R^{\mathfrak{X}})_{R \in \sigma})$ and a set \tilde{A} of partial permutations of X that, in fact, are partial automorphisms of \mathfrak{X} . The pair $(\mathfrak{X}, \tilde{A})$ is called an *extension problem*. A pair $(\mathfrak{Y}, \{\hat{a}: a \in \tilde{A}\})$ is a *solution* for the extension problem $(\mathfrak{X}, \tilde{A})$ if \mathfrak{Y} is an extension of \mathfrak{X} , \hat{a} is an automorphism of \mathfrak{Y} , $a \subseteq \hat{a}$ and $\hat{a}^{-1} = \widehat{a^{-1}}$ for every $a \in \tilde{A}$. If the automorphisms \hat{a} are clear from context, we occasionally refer to the structure \mathfrak{Y} itself as a solution for the extension problem $(\mathfrak{X}, \tilde{A})$.

We shall apply the concepts of the preceding subsection. Let G be a not necessarily finite A -generated group extending the graph \mathcal{P} given by the data (X, \tilde{A}) according to Definition 2.3. In order to equip all of $Y := X_{\mathcal{P}}G$ with a relational structure such that every $g \in G$ is not only a bijection on Y but an automorphism of \mathfrak{Y} it is natural to add to the interpretations of the various symbols R on \mathfrak{X} whatever is required to make g an automorphism on all of \mathfrak{Y} . For every $R \in \sigma$ set

$$R^{\mathfrak{Y}} := \bigcup \{R^{\mathfrak{X}}g: g \in G\}$$

where $R^{\mathfrak{X}}g = \{(x_1g, \dots, x_rg): (x_1, \dots, x_r) \in R^{\mathfrak{X}}\}$. In this way, $\mathfrak{Y} =: \mathfrak{X}_{\mathcal{P}}G$ becomes a σ -structure and all $g \in G$ act as automorphisms of \mathfrak{Y} . The structure \mathfrak{X} is clearly a weak substructure of \mathfrak{Y} . We give a sufficient condition on G for \mathfrak{X} to be an induced substructure of \mathfrak{Y} , that is, for \mathfrak{Y} being a solution for the extension problem $(\mathfrak{X}, \tilde{A})$.

Lemma 2.6. *Let \mathfrak{X} be a finite σ -structure on base set X and \tilde{A} a set of partial automorphisms of \mathfrak{X} , encoded in the A -graph \mathcal{P} on the vertex set X . Let G be an A -generated group which extends \mathcal{P} and let $M = \mathcal{T}(\mathcal{P})$. Suppose that G satisfies the following condition: for each $s \leq |X|$ and each collection of words $w_1, \dots, w_s \in \tilde{A}^*$ for which $[w_1]_G = \dots = [w_s]_G$ there*

exists a word $w \in \tilde{A}^*$ such that $[w_1]_M, \dots, [w_s]_M \leq [w]_M$. Then \mathfrak{X} is an induced substructure of $\mathfrak{X}_{\mathcal{P}}G$.

Proof. Let $x_1, \dots, x_r \in X$ and suppose that $(x_1, \dots, x_r) \in R^{\mathfrak{X}_{\mathcal{P}}G}$ for some $R \in \sigma$. Then there are $y_1, \dots, y_r \in X$ and $g \in G$ such that $(y_1, \dots, y_r) \in R^{\mathfrak{X}}$ and $x_i = y_i g$ for $i = 1, \dots, r$. By Lemma 2.2 there exist words $w_1, \dots, w_r \in \tilde{A}^*$ such that $y_i \cdot [w_i]_M = x_i$ and $[w_i]_G = g$. If $r > |X|$ then the r -tuple (y_1, \dots, y_r) nevertheless contains at most $|X|$ distinct elements. In any case, we may assume that the collection of words w_1, \dots, w_r contains at most $|X|$ distinct words. By the assumption of the Lemma there exists a word $w \in \tilde{A}^*$ for which $[w_i]_M \leq [w]_M$ for all i . This implies $y_i \cdot [w]_M = x_i$ for all i . Since $[w]_M$ is a partial automorphism of \mathfrak{X} it follows that $(x_1, \dots, x_s) \in R^{\mathfrak{X}}$. \square

Remark 2.7. If in Lemma 2.6 the largest arity r of σ is smaller than the cardinality $|X|$ then the bound $|X|$ may be replaced by r .

Readers familiar with semigroup theory will observe that the condition on G , in relation to the inverse monoid M , formulated in Lemma 2.6 just says that G witnesses the group-pointlike sets of M of order up to $|X|$.

2.5.3. Classes defined by forbidden homomorphisms and free extensions of relational structures. Let σ be a finite relational signature and Σ a finite set of finite σ -structures; let $\mathbf{Excl}(\Sigma)$ be the class of all σ -structures \mathfrak{S} which do not admit a homomorphism $\mathfrak{U} \rightarrow \mathfrak{S}$ for any $\mathfrak{U} \in \Sigma$. A structure \mathfrak{S} which does not admit such a homomorphism is said to be Σ -free or to *avoid* (all members of) Σ . The members of Σ are called the structures *forbidden for the class* $\mathbf{Excl}(\Sigma)$. A class \mathbf{C} of σ -structures is said to have the *extension property for partial automorphisms* (EPPA for short) if every extension problem $(\mathfrak{S}, \tilde{A})$ for finite $\mathfrak{S} \in \mathbf{C}$ that admits a solution in \mathbf{C} also admits a finite solution in \mathbf{C} . Given a finite σ -structure $\mathfrak{X} = (X; (R^{\mathfrak{X}})_{R \in \sigma})$ and a set of partial automorphisms \tilde{A} of \mathfrak{X} encoded in the finite A -graph \mathcal{P} , we can extend \mathfrak{X} by means of the free A -generated group F , as explained in Section 2.5.2 (F satisfies the condition formulated in Lemma 2.6). The σ -structure on $X_{\mathcal{P}}F$ so obtained is called the *free extension* $\mathfrak{X}_{\mathcal{P}}F$ (subject to \mathcal{P}) of \mathfrak{X} . Suppose that $(\mathfrak{X}, \tilde{A})$ has some solution $(\mathfrak{Y}, \{\hat{a}: a \in \tilde{A}\})$; then $G := \langle \hat{a}: a \in \tilde{A} \rangle$, is an A -generated subgroup of the automorphism group $\text{Aut}(\mathfrak{Y})$ of \mathfrak{Y} . If \mathfrak{Y} is \mathfrak{U} -free for some finite σ -structure \mathfrak{U} then so is $[\mathfrak{X}]G$, the substructure of \mathfrak{Y} induced on the orbit of \mathfrak{X} under the action of G . It follows from principles of universal algebra that the identity mapping $X \rightarrow X$, $x \mapsto x$ extends to surjective homomorphisms of \tilde{A} -sets $X_{\mathcal{P}}F \twoheadrightarrow X_{\mathcal{P}}G \twoheadrightarrow [X]G$ (the latter being the base set of $[\mathfrak{X}]G$). From the definition of the σ -structures on $\mathfrak{X}_{\mathcal{P}}F$ and $\mathfrak{X}_{\mathcal{P}}G$, respectively, it follows that we even have surjective homomorphisms of the relational structures $\mathfrak{X}_{\mathcal{P}}F \twoheadrightarrow \mathfrak{X}_{\mathcal{P}}G \twoheadrightarrow [\mathfrak{X}]G$. Summing up, this means that, if \mathfrak{X} has some \mathfrak{U} -free extension then the free extension $\mathfrak{X}_{\mathcal{P}}F$ is also \mathfrak{U} -free.

2.6. Stallings graphs and Schreier graphs. We introduce the Stallings and Schreier graphs assigned to a finitely generated subgroup H of the free A -generated group F ; we shall also consider the appropriate concepts for cosets Hg .

Let H be a finitely generated subgroup of F and let w_1, \dots, w_n be generating elements, realised as reduced words over \tilde{A} . Consider the disjoint union of cycle graphs $\mathcal{C}_1, \dots, \mathcal{C}_n$, where \mathcal{C}_i is the graph spanned by a closed simple path of length $|w_i|$ labelled w_i at a base vertex ι_i . Next factor this disjoint union of labelled cycle graphs \mathcal{C}_i by the smallest A -graph congruence that identifies all base vertices ι_i . The equivalence class containing all vertices ι_i then is a distinguished vertex ι of the quotient graph. This pointed graph with distinguished vertex ι is called the *Stallings graph* \mathcal{S}_H of the group H . It depends just on H , not on the choice of the generating elements w_1, \dots, w_n (see Proposition 2.8). It can be obtained in a nice manner from the cycle graphs \mathcal{C}_i by a procedure as follows: first identify all base vertices ι_1, \dots, ι_n with each other and denote the resulting vertex as ι and then apply Stallings foldings until the result is an A -graph; for more details the reader may consult [21].

The essential feature of the Stallings graph \mathcal{S}_H is this: the group H is comprised exactly of the (reduced) words labelling reduced paths in \mathcal{S}_H that are closed at the vertex ι . The Stallings graph \mathcal{S}_H is an invariant of the group H in the strong sense of Proposition 2.8 below. In order to formulate it we need another concept: for a finite connected A -graph \mathcal{P} with distinguished vertex ι let $L(\iota, \mathcal{P})$ denote the set of all words over \tilde{A} that are labels of reduced paths in \mathcal{P} that are closed at ι . All words in $L(\iota, \mathcal{P})$ are reduced, and it is easy to see that $L(\iota, \mathcal{P})$ is a finitely generated subgroup of F . By construction, $L(\iota, \mathcal{S}_H) = H$. The uniqueness of the Stallings graph \mathcal{S}_H is characterised by the following proposition.

Proposition 2.8 ([21], Theorem 5.2). *Let H be a finitely generated subgroup of F and \mathcal{P} a finite connected A -graph with distinguished vertex ι such that every vertex of \mathcal{P} except possibly ι has degree at least 2. If $L(\iota, \mathcal{P}) = H$ then there is a unique isomorphism of pointed graphs $\mathcal{S}_H \rightarrow \mathcal{P}$.*

We can generalise this concept to cosets Hg , already discussed in Section 4.17 in the monograph [31]. Let H be given, \mathcal{S}_H be its Stallings graph and let $g \in F$ be represented as a reduced word $g = a_1 \cdots a_n \in \tilde{A}^*$. Take the graph \mathcal{L} spanned by a path $p_0 \xrightarrow{e_1} p_1 \xrightarrow{e_2} \cdots \xrightarrow{e_n} p_n$ where $n = |g|$ is the length of the word g . We label the edge e_i by the letter a_i so that the label of the path $e_1 \cdots e_n$ is just g (with inverse edges labelled accordingly). We form the disjoint union of \mathcal{S}_H and \mathcal{L} and factor by the smallest A -graph congruence that identifies the distinguished vertex ι of \mathcal{S}_H and the initial vertex p_0 of \mathcal{L} . In other words, we identify the vertex p_0 with the vertex ι of \mathcal{S}_H and perform Stallings foldings until the result is an A -graph. Denoting the vertex corresponding to the terminal vertex p_n of the path in

this graph by τ , we obtain a resulting, 2-pointed graph \mathcal{S}_{Hg} , with ι and τ as distinguished vertices. We call ι the *initial* vertex and τ the *terminal* vertex of \mathcal{S}_{Hg} , and refer to the 2-pointed graph \mathcal{S}_{Hg} as the *Stallings graph of the coset Hg* . This term is justified since Hg is comprised of all words that are labels of reduced paths running in \mathcal{S}_{Hg} with initial vertex ι and terminal vertex τ . It may happen that $\iota = \tau$; this is exactly the case when $H = Hg$, that is, when $g \in H$. For Stallings graphs of cosets a characterisation analogous to Proposition 2.8 holds. For a finite connected A -graph \mathcal{P} with two distinguished vertices ι and τ let $L(\iota, \mathcal{P}, \tau)$ be the set of all words that are labels of reduced paths $\pi: \iota \rightarrow \tau$ running in \mathcal{P} . Every word of $L(\iota, \mathcal{P}, \tau)$ is reduced and $L(\iota, \mathcal{P}, \tau)$ is a coset Hg with H finitely generated and g a word labelling a path $\iota \rightarrow \tau$ in \mathcal{P} . The proof of the next proposition is completely analogous to the one of Proposition 2.8.

Proposition 2.9. *Let H be a finitely generated subgroup of F and $g \in F$; let \mathcal{P} be a finite connected A -graph with distinguished vertices ι and τ such that all vertices of \mathcal{P} except possibly ι and τ have degree at least 2. If $L(\iota, \mathcal{P}, \tau) = Hg$ then there is a unique isomorphism of 2-pointed graphs $\mathcal{S}_{Hg} \xrightarrow{\sim} \mathcal{P}$.*

Moreover, in a Stallings graph \mathcal{S}_{Hg} with $\iota \neq \tau$ we may swap the rôles of ι and τ , that is, declare ι to be the terminal vertex and τ to be the initial vertex. By doing so we get the Stallings graph of the coset $(g^{-1}H)g^{-1}$ where $g^{-1}H$ is the conjugate group $g^{-1}Hg$ (note that the right coset $(g^{-1}H)g^{-1}$ coincides with the left coset $g^{-1}H$). Indeed, the labels of reduced paths in \mathcal{S}_{Hg} starting at τ and ending at ι are exactly the inverses of the labels of the reduced paths starting at ι and ending at τ .

Stallings graphs \mathcal{S}_{Hg} of cosets Hg will play a crucial rôle in the proofs of the implications Herwig–Lascar \implies Ribes–Zalesskii and Herwig–Lascar \implies Ash. The free completion $\mathcal{S}_{Hg}F$ of \mathcal{S}_{Hg} , which we call the *Schreier graph* of the coset Hg , will also be important in this context. We shall denote it as \mathbf{S}_{Hg} and consider it also as a 2-pointed graph with initial vertex ι and terminal vertex τ . The behaviour of \mathbf{S}_{Hg} , compared to that of \mathcal{S}_{Hg} , can be summarised as follows: for a *reduced* word $w \in \tilde{A}^*$, the equality $\iota \cdot w = \tau$ holds in \mathcal{S}_{Hg} if and only if $w \in Hg$, while for *any* word $w \in \tilde{A}^*$ the equality $\iota w = \tau$ holds in \mathbf{S}_{Hg} if and only if $\text{red}(w) \in Hg$. In contrast to the Schreier graph \mathbf{S}_{Hg} , the Stallings graph \mathcal{S}_{Hg} is not complete (unless H is of finite index [21, Proposition 8.3]). It may happen that for a word w for which $\text{red}(w)$ belongs to Hg , $\iota \cdot w$ is not defined in \mathcal{S}_{Hg} .

Usually, the Schreier graph of a subgroup H of F is defined to be the A -graph with vertex set $\{Hg: g \in F\}$, the set of all right cosets with respect to H , and edges $e: Hg \xrightarrow{a} Hg[a]_F$, labelled a , for all $g \in F$ and $a \in \tilde{A}$. It is therefore sometimes called the *coset graph* or *Schreier coset graph* with respect to H . It is known from [21] that \mathcal{S}_HF and the latter graph are isomorphic as pointed A -graphs, via a unique isomorphism mapping $\iota \mapsto H$, see [21]. The Stallings graph \mathcal{S}_H is the subgraph of the Schreier graph \mathbf{S}_H spanned by all reduced paths closed at the distinguished vertex ι . The latter

could be used as an alternative definition of the Stallings graph based on the alternative (usual) definition of the Schreier graph as coset graph with respect to H . For our purposes the view of the Schreier graph as $\mathcal{S}_H = \mathcal{S}_H F$ seems to be more appropriate.

We close this section by an easy and well-known, but important observation. Let M be a finite A -generated inverse monoid; for $m \in M$ we let

$$\varphi(m) := \{[w]_F : w \in \tilde{A}^*, [w]_M = m\}. \quad (2.3)$$

Obviously, φ is a mapping from M to the powerset 2^F of F , which, in semigroup-theoretic terms, is the *canonical relational morphism* $M \rightarrow F$.

Proposition 2.10. *Let M be a finite A -generated inverse monoid and $m \in M$; then $\varphi(m) = Hg$ for some finitely generated subgroup H of F and some $g \in F$.*

Proof. Let $m \in M$ and recall from Section 2.4 the Schützenberger graph \mathcal{M}_m with respect to m . Its set of vertices is R_m , the \mathcal{R} -class of m , and for $p, r \in R_m$ and $a \in \tilde{A}$ there is an edge $p \rightarrow r$ labelled a if and only if $r = p[a]_M$. Let H be the set of all reduced words labelling (necessarily reduced) paths in \mathcal{M}_m that are closed at the vertex mm^{-1} . Since \mathcal{M}_m is finite, H is a finitely generated subgroup of F and, by Proposition 2.8, \mathcal{S}_H is a subgraph of \mathcal{M}_m with distinguished vertex $\iota = mm^{-1}$. Moreover, \mathcal{S}_H contains every edge of \mathcal{M}_m that is contained in a non-trivial cycle subgraph of \mathcal{M}_m . In other words, \mathcal{M}_m possibly differs from \mathcal{S}_H just by finite trees attached to \mathcal{S}_H . Let g be a reduced word labelling a (necessarily reduced) path $mm^{-1} \rightarrow m$. Then by Proposition 2.9, the Stallings graph \mathcal{S}_{Hg} of the coset Hg is also a subgraph of \mathcal{M}_m with distinguished vertices $\iota = mm^{-1}$ and $\tau = m$, respectively. Again, \mathcal{M}_m and \mathcal{S}_{Hg} differ at most by finite trees attached to \mathcal{S}_{Hg} . Let $u \in \tilde{A}^*$ be such that $[u]_F \in \varphi(m)$, that is, $[u]_M = m$; then u labels a path $mm^{-1} \rightarrow m$ in \mathcal{M}_m because $m = mm^{-1}m = mm^{-1}[u]_M$. The reduced form $\text{red}(u)$ still labels a path $mm^{-1} \rightarrow m$, but this path now runs in the subgraph \mathcal{S}_{Hg} (by the aforementioned difference between \mathcal{M}_m and \mathcal{S}_{Hg}), hence $[u]_F = \text{red}(u) \in Hg$. Let conversely $f \in Hg$ be represented as a reduced word over \tilde{A} . Then f labels a path $mm^{-1} \rightarrow m$ in \mathcal{S}_{Hg} and hence also in \mathcal{M}_m . According to Lemma 2.1, $m \leq [f]_M$. Let again u be a word with $[u]_M = m$. Then $[uu^{-1}f]_M = m$ and $[uu^{-1}f]_F = f$, hence $f \in \varphi(m)$ as required. \square

3. THEOREMS REVISITED

We formulate the three theorems in question.

Theorem 3.1 (Ribes–Zalesskii [32]). *For every positive integer n , the product $H_1 \cdots H_n$ of every n -tuple of finitely generated subgroups H_1, \dots, H_n of the A -generated free group F is closed in the profinite topology of F .*

Recall that a basis for the open sets of the profinite topology of F is given by the set of all (left or right) cosets Hg with $g \in F$ and H a finite index

subgroup of F . In fact, it suffices to consider cosets Ng with N a normal subgroup of finite index. From this it follows that a subset $W \subseteq F$ is closed if and only if for every element $g \notin W$ there exists a finite index normal subgroup N of F such that $Ng \cap W = \emptyset$.

Theorem 3.2 (Herwig–Lascar [19]). *Let σ be a finite relational signature; every class \mathbf{C} of σ -structures that is defined in terms of forbidden homomorphisms of finitely many finite σ -structures enjoys the extension property for partial automorphisms.*

In order to formulate Ash’s Theorem we need some further terminology. Given a graph \mathcal{E} , an A -generated group G and an A -generated inverse monoid M . We say that a word labelling $v: K(\mathcal{E}) \rightarrow \tilde{A}^*$ *commutes* over the group G (is G -commuting, for short) if, for every closed path π in \mathcal{E} the G -value of the label of π is the identity element: $[\ell(\pi)]_G = 1_G$; we say two labellings $u, v: K(\mathcal{E}) \rightarrow \tilde{A}^*$ are M -related if $[u(e)]_M = [v(e)]_M$ for every edge $e \in K(\mathcal{E})$. The following is essentially the same as Theorem 5.1 in [6], but slightly reformulated. However, this is what was literally proved in Sections 5 – 7 in [6]. The original, different formulation of Ash’s Theorem 5.1 came from the originally intended purpose of this theorem.

Theorem 3.3 (Ash [6], inverse monoid version). *For every finite A -generated inverse monoid M and every finite graph \mathcal{E} there exists a finite A -generated group G such that every G -commuting word labelling of \mathcal{E} admits an M -related relabelling that commutes over the free group F .*

If a word labelling of \mathcal{E} does not admit an M -related and F -commuting relabelling, then, according to the theorem, this labelling is not G -commuting. Hence G may be interpreted as a finite device for eliminating labellings that are “bad” with respect to having M -related relabellings that commute over F .

3.1. Herwig–Lascar implies Ribes–Zalesskii. Let $n \geq 2$; we consider relational structures over the signature $\sigma_n = \{R_1, \dots, R_n\}$ consisting of n binary relational symbols. For elements y_1, y_2 in such a structure $\mathfrak{Y} = (Y; R_1^{\mathfrak{Y}}, \dots, R_n^{\mathfrak{Y}})$ in question we shall often write $y_1 R_i y_2$ instead of $(y_1, y_2) \in R_i^{\mathfrak{Y}}$ (in particular, we often shall omit the superscript \mathfrak{Y} in $R_i^{\mathfrak{Y}}$ whenever there is no danger of confusion). Such structures can be seen as *edge-coloured directed graphs*, though they are not graphs in the sense of Section 2.3 and we do not use the term graph in this context in order to avoid confusion.

We consider the class $\mathbf{CCF}_n := \mathbf{Excl}(\mathfrak{C}_n)$ of all σ_n -structures \mathfrak{C} that admit no homomorphism $\mathfrak{C}_n \rightarrow \mathfrak{C}$ from the σ_n -structure \mathfrak{C}_n consisting of n distinct elements c_1, \dots, c_n with

$$R_i^{\mathfrak{C}_n} = \{(c_i, c_{i+1})\} \text{ for } i = 1, \dots, n-1 \text{ and } R_n^{\mathfrak{C}_n} = \{(c_n, c_1)\}.$$

(The structure \mathfrak{C}_n may be seen as the directed cycle graph of length n whose edges have colours $1, \dots, n$.) According to Theorem 3.2, \mathbf{CCF}_n has

the *extension property for partial automorphisms*: every extension problem $(\mathfrak{S}, \tilde{A})$ for finite \mathfrak{S} in \mathbf{CCF}_n that admits a solution in \mathbf{CCF}_n also has a finite solution in \mathbf{CCF}_n .

The following proof of the implication Herwig–Lascar \implies Ribes–Zalesskii is perhaps not the shortest one, but it tries to prepare the ground for a direct proof of the implication Herwig–Lascar \implies Ash, which will be the content of Section 3.4. The proof is inspired by the proof of the Ribes–Zalesskii Theorem given in [19] but, in a sense, embeds those ideas of [19] into the environment of Stallings and Schreier graphs (which seem to be the natural framework for this topic). Crucial will be the following result.

Proposition 3.4. *Let $n \geq 2$, let H_1, \dots, H_n be finitely generated subgroups of F and $g_1, \dots, g_n \in F$. Then there exists a finite index normal subgroup N of F such that: if $1_F \notin H_1 g_1 \cdots H_n g_n$ then $1_F \notin N H_1 g_1 \cdots N H_n g_n$.*

Proof. We derive this result from the extension property for partial automorphisms for the class \mathbf{CCF}_n . So, let H_1, \dots, H_n be finitely generated subgroups of F and $g_1, \dots, g_n \in F$ and, for each i , let $\mathcal{S}_i \subseteq \mathfrak{S}_i = \mathfrak{S}_i F$ be the Stallings and Schreier graphs of the coset $H_i g_i$ (see Section 2.6) with distinguished vertices ι_i and τ_i , respectively.

Consider the disjoint unions

$$\mathcal{S}_1 \sqcup \cdots \sqcup \mathcal{S}_n \subseteq \mathfrak{S}_1 \sqcup \cdots \sqcup \mathfrak{S}_n.$$

The sets of vertices of these graphs will be the base sets for the structures to be defined below. The labelled edges shall indicate (partial) automorphisms of the structures in question. We first note that the free group F acts by bijections on the vertex set of $\mathfrak{S} := \mathfrak{S}_1 \sqcup \cdots \sqcup \mathfrak{S}_n$ on the right via $\rho \mapsto \rho w$ for every vertex ρ and every $w \in F$, and for each individual i the vertex set of \mathfrak{S}_i is invariant under this action.

We define binary relations R_1, \dots, R_n on the vertex set $S = V(\mathfrak{S})$ of \mathfrak{S} in order to produce a structure in \mathbf{CCF}_n provided that $1_F \notin H_1 g_1 \cdots H_n g_n$. We set

$$\tau_1 R_1 \iota_2, \tau_2 R_2 \iota_3, \dots, \tau_{n-1} R_{n-1} \iota_n, \tau_n R_n \iota_1 \quad (3.1)$$

and add to this whatever is required so that the aforementioned action $\rho \mapsto \rho w$ by F is an action by automorphisms of the structure. More precisely:

$$R_i = \{(\tau_i w, \iota_{i+1} w) : w \in F\} \text{ for } i = 1, \dots, n-1$$

and

$$R_n = \{(\tau_n w, \iota_1 w) : w \in F\}.$$

We first show that, if $1_F \notin H_1 g_1 \cdots H_n g_n$ then $\mathfrak{S} := (S; R_1, \dots, R_n)$ is a member of \mathbf{CCF}_n . Suppose that $\mathfrak{S} = (S; R_1, \dots, R_n) \notin \mathbf{CCF}_n$. This means that there is a “coloured cycle”: elements ν_1, \dots, ν_n such that

$$\nu_1 R_1 \nu_2 R_2 \nu_3 \cdots \nu_{n-1} R_{n-1} \nu_n R_n \nu_1. \quad (3.2)$$

We note that $\nu_i \in V(\mathfrak{S}_i)$ for every i , because every set $V(\mathfrak{S}_i)$ is invariant under the action of F . By definition of the relations R_i , there are

$w_1, w_2, \dots, w_n \in F$ such that

$$\begin{aligned} \nu_1 &= \tau_1 w_1 \\ \iota_2 w_1 &= \nu_2 = \tau_2 w_2 \\ &\vdots \\ \iota_n w_{n-1} &= \nu_n = \tau_n w_n \\ \iota_1 w_n &= \nu_1 \end{aligned}$$

This situation is depicted in Figure 1. Recall the rôles of the vertices ι_i and




FIGURE 1. The structure \mathfrak{S}

τ_i in the Schreier graph of the coset $H_i g_i$: from the equality $\iota_i w_{i-1} w_i^{-1} = \tau_i = \iota_i g_i$ (which holds in \mathfrak{S}_i for $i = 2, \dots, n$) and $\iota_1 w_n w_1^{-1} = \tau_1 = \iota_1 g_1$ (which holds in \mathfrak{S}_1) it follows that

$$w_n w_1^{-1} \in H_1 g_1, w_1 w_2^{-1} \in H_2 g_2, \dots, w_{n-1} w_n^{-1} \in H_n g_n$$

and so

$$1_F = w_n w_1^{-1} w_1 w_2^{-1} \dots w_{n-1} w_n^{-1} \in H_1 g_1 H_2 g_2 \dots H_n g_n.$$

Altogether, if $1_F \notin H_1 g_1 \dots H_n g_n$ then $\mathfrak{S} = (\mathcal{S}; R_1, \dots, R_n) \in \mathbf{CCF}_n$.

Now consider the substructure \mathfrak{V} of \mathfrak{S} induced on the finite subset $V = V(\mathfrak{S})$ where

$$\mathcal{S} := \mathcal{S}_1 \sqcup \dots \sqcup \mathcal{S}_n. \quad (3.3)$$

First note that (3.1) holds in \mathfrak{V} . The A -graph structure on \mathcal{S} indicates partial automorphisms of the σ_n -structure \mathfrak{V} . Since \mathfrak{S} is an extension in \mathbf{CCF}_n in which all these partial automorphisms extend to total automorphisms, there exists a finite such extension, say $\overline{\mathfrak{V}} = (\overline{V}; R_1, \dots, R_n) \in \mathbf{CCF}_n$. Recall that

we intend to show that there exists a finite index normal subgroup N of F such that $1_F \notin NH_1g_1 \cdots NH_ng_n$.

We obtain the normal subgroup N from the finite A -generated group \tilde{G} of automorphisms of $\tilde{\mathfrak{V}}$ that extend the partial automorphisms given by \tilde{A} on \mathfrak{V} , by putting $F/N := G$. The free group F acts on $\tilde{\mathfrak{V}}$ on the right by automorphisms via $\rho w := \rho[w]_G$. For $\rho \in \mathfrak{V}$ and $w \in F$, if $\rho \cdot w$ is defined in \mathfrak{V} (where $\rho \mapsto \rho \cdot w$ denotes the partial bijection on $V = V(\mathcal{S})$ induced by w) then it has the same value in $\tilde{\mathfrak{V}}$, that is $\rho \cdot w = \rho[w]_G = \rho w$.

We claim that $1_F \notin NH_1g_1 \cdots NH_ng_n$. Suppose towards a contradiction this were not the case. Then, for every $i = 1, \dots, n$ there is $k_i \in NH_ig_i$ such that $1_F = k_1k_2 \cdots k_n$. We note that $\iota_i k_i = \tau_i$ for every i : indeed, $k_i = n_i h_i g_i$ for some $n_i \in N$ and h_i in H_i . Now $\iota_i n_i = \iota_i$ since $[n_i]_G = 1_G$. Moreover, $\iota_i \cdot h_i = \iota_i$ and $\iota_i \cdot g_i = \tau_i$ hold in \mathfrak{V} . By use of (3.1), which holds in \mathfrak{V} , the following hold in $\tilde{\mathfrak{V}}$:

$$\begin{aligned} & \tau_1 k_2 \cdots k_n R_1 \iota_2 k_2 \cdots k_n \\ &= \tau_2 k_3 \cdots k_n R_2 \iota_3 k_3 \cdots k_n \\ & \quad \vdots \\ &= \tau_{n-1} k_n R_{n-1} \iota_n k_n \\ &= \tau_n R_n \iota_1. \end{aligned}$$

Since $\iota_1 = \iota_1 k_1 \cdots k_n = \tau_1 k_2 \cdots k_n$ this leads to a “coloured cycle” as in (3.2), a contradiction to the fact that $\tilde{\mathfrak{V}} \in \mathbf{CCF}_n$. \square

The Ribes–Zalesskii Theorem is an immediate consequence: let H_1, \dots, H_n be finitely generated subgroups of F and $g \in F$. Suppose that $g \notin H_1 \cdots H_n$; this is equivalent to $1_F \notin H_1 \cdots H_n g^{-1}$. By Proposition 3.4 this implies $1_F \notin NH_1 \cdots NH_n g^{-1} = NH_1 \cdots H_n g^{-1}$ for some finite index normal subgroup N of F . Hence the set $NH_1 \cdots H_n g^{-1}$ is a union of finitely many cosets Nf_1, \dots, Nf_s (since N is finite index) for which $Nf_i \neq N$ for all i . It follows that $N \cap H_1 \cdots H_n g^{-1} \subseteq N \cap NH_1 \cdots H_n g^{-1} = \emptyset$. Multiplication by g on the right gives $Ng \cap H_1 \cdots H_n = \emptyset$.

3.2. Ribes–Zalesskii implies Ash. The following is a consequence of the Ribes–Zalesskii Theorem:

Theorem 3.5. *For every finite A -generated inverse monoid M and every positive integer n there exists a finite A -generated group G such that for every n -tuple v_1, \dots, v_n of words in \tilde{A}^* for which $[v_1 \cdots v_n]_G = 1_G$ there exists an n -tuple of words u_1, \dots, u_n in \tilde{A}^* such that*

- (1) $[u_i]_M = [v_i]_M$ for all i , and
- (2) $[u_1 \cdots u_n]_F = 1_F$.

Proof. We derive this claim from the Theorem of Ribes–Zalesskii. We note that for every finitely generated subgroup H of F , every conjugate subgroup

${}^wH := wHw^{-1}$ is also finitely generated. Hence every subset of F of the form

$$H_1w_1 \cdot H_2w_2 \cdots H_nw_n$$

for the H_i finitely generated subgroups of F and $w_i \in F$ is also closed. Indeed, one has

$$H_1w_1 \cdot H_2w_2 \cdots H_nw_n = H_1 \cdot {}^{w_1}H_2 \cdot {}^{w_1w_2}H_3 \cdots {}^{w_1 \cdots w_{n-1}}H_nw_1w_2 \cdots w_n,$$

so the right hand side is a product of n finitely generated subgroups of F multiplied by the element $w_1 \cdots w_n$. Since right multiplication $x \mapsto xw$ for some fixed element $w \in F$ is a homeomorphism of F with respect to the profinite topology, the latter set is also closed.

Let $\varphi: M \rightarrow 2^F$ be the mapping defined by

$$\varphi(m) := \{[w]_F : w \in \tilde{A}^*, [w]_M = m\}.$$

Proposition 2.10 tells us that for every $m \in M$, $\varphi(m) = Hw$ for some finitely generated subgroup H of F and some $w \in F$.

For every n -tuple (m_1, m_2, \dots, m_n) of elements $m_i \in M$ consider the corresponding set

$$\varphi(m_1) \cdot \varphi(m_2) \cdots \varphi(m_n) = H_1w_1 \cdot H_2w_2 \cdots H_nw_n,$$

which is a closed subset of F . Such a set may or may not contain the identity element 1_F . For every such set $H_1w_1 \cdot H_2w_2 \cdots H_nw_n$ that does not contain 1_F there exists a finite index normal subgroup $N(m_1, \dots, m_n)$ of F such that $N(m_1, \dots, m_n) \cap H_1w_1 \cdot H_2w_2 \cdots H_nw_n = \emptyset$. The intersection of all such (finitely many) normal subgroups is again a finite index normal subgroup N of F . By definition, N satisfies the following. For every n -tuple $(m_1, \dots, m_n) \in M^n$, if $1_F \notin \varphi(m_1) \cdots \varphi(m_n)$, then $N \cap \varphi(m_1) \cdots \varphi(m_n) = \emptyset$. The finite group $G := F/N$ fulfils the requirements of the theorem. Indeed, let $v_1, \dots, v_n \in \tilde{A}^*$ be such that $[v_1 \cdots v_n]_G = 1_G$, which means that $[v_1 \cdots v_n]_F \in N$. We let $m_i := [v_i]_M$ for $i = 1, \dots, n$. Since $[v_i]_F \in \varphi(m_i)$, we also have

$$[v_1 \cdots v_n]_F = [v_1]_F \cdots [v_n]_F \in \varphi(m_1) \cdots \varphi(m_n).$$

So in particular $N \cap \varphi(m_1) \cdots \varphi(m_n) \neq \emptyset$, from which it follows by construction of N that $1_F \in \varphi(m_1) \cdots \varphi(m_n)$. Consequently, $1_F = g_1 \cdots g_n$ for some $g_i \in \varphi(m_i)$. By definition of φ there exist words u_1, \dots, u_n such that $[u_i]_F = g_i$ and $[u_i]_M = m_i = [v_i]_M$ for every i and

$$1_F = g_1 \cdots g_n = [u_1 \cdots u_n]_F.$$

□

It should be mentioned that, conversely, Theorem 3.5 implies the Ribes–Zalesskii Theorem. Indeed let H_1, \dots, H_n be finitely generated subgroups of F and let $g \in F$ (considered as a reduced word over \tilde{A}). Take the graph \mathcal{S} in (3.3) with \mathcal{S}_i the Stallings graph of H_i for $i = 1, \dots, n-1$ and \mathcal{S}_n the Stallings

graph of the coset $H_n g^{-1}$, and let $M = \mathcal{T}(\mathcal{S})$.² We have that $g \in H_1 \cdots H_n$ if and only if $1_F \in H_1 \cdots H_n g^{-1}$ and it is easy to see that the latter holds if and only if there are words $u_1, \dots, u_n \in \tilde{A}^*$ such that $[u_1 \cdots u_n]_F = 1_F$ and $\iota_i \cdot u_i = \iota_i$ for $i = 1, \dots, n-1$ and $\iota_n \cdot u_n = \iota_n \cdot g^{-1} = \tau_n$ holds in \mathcal{S} . Let G be a finite A -generated group satisfying Theorem 3.5 for the inverse monoid M . Let $\varphi: F \rightarrow G$ be the canonical morphism and suppose that $\varphi(g) \in \varphi(H_1 \cdots H_n)$. Then there exist (reduced) words h_1, \dots, h_n with $h_i \in H_i$ for all i such that $\varphi(g) = \varphi(h_1 \cdots h_n)$, that is, $[g]_G = [h_1 \cdots h_n]_G$ so that $[h_1 \cdots h_n g^{-1}]_G = 1_G$. By Theorem 3.5 there exist words u_1, \dots, u_n such that $[u_1 \cdots u_n]_F = 1_F$, $[u_i]_M = [h_i]_M$ (for $i = 1, \dots, n-1$) and $[u_n]_M = [h_n g^{-1}]_M$. The latter condition implies $\iota_i \cdot u_i = \iota_i$ (for $i = 1, \dots, n-1$) and $\iota_n \cdot u_n = \iota_n \cdot g^{-1} = \tau_n$, which in turn gives $g \in H_1 \cdots H_n$. Consequently, if $g \notin H_1 \cdots H_n$, then g and $H_1 \cdots H_n$ are separated by the morphism φ .

Observe that Theorem 3.5 is Theorem 3.3 for the special case of cycle graphs. We have thus proved the equivalence of Ash's Theorem for cycle graphs with the Theorem of Ribes–Zaleskii (a more long-winded proof for this equivalence may be found in [37].) We now derive Theorem 3.3 from Theorem 3.5. The following concept will be convenient for the proof. It is essentially the same as that of an “associated tree” in [6, Section 7].

Definition 3.6 (critical group). Let M be a finite A -generated inverse monoid, \mathcal{E} a finite connected graph without loop edges and G a finite A -generated group extending M . We say that G is *critical for the pair* (M, \mathcal{E}) if the following condition is satisfied. For every G -commuting labelling $v: K(\mathcal{E}) \rightarrow \tilde{A}^*$ there exists a finite tree \mathcal{T} of diameter at most $|V(\mathcal{E})| - 1$ such that $V(\mathcal{E}) \subseteq V(\mathcal{T})$, and a labelling $w: K(\mathcal{T}) \rightarrow \tilde{A}^*$ such that for every edge $e \in K(\mathcal{E})$, if π_e denotes the (unique) reduced path $\alpha e \rightarrow \omega e$ in \mathcal{T} then $[v(e)]_M \leq [w(\pi_e)]_M$.

Note that every word labelling of a finite tree is F -commuting. A group G critical for (M, \mathcal{E}) guarantees that every labelling of \mathcal{E} that is acknowledged by G to commute admits an M -related F -commuting relabelling, as we will see immediately after the statement of the following key lemma of this subsection.

Lemma 3.7 (key lemma). *For every finite A -generated inverse monoid M and every finite connected graph \mathcal{E} without loop edges there exists a finite A -generated group G which is critical for the pair (M, \mathcal{E}) .*

We first show how this lemma implies Ash's Theorem.

Proof of Theorem 3.3 by use of Lemma 3.7. Let M be a finite A -generated inverse monoid and let \mathcal{E} be a finite connected graph. Let \mathcal{D} be the graph obtained from \mathcal{E} by removing all loop edges (if any). Suppose that G is critical for the pair (M, \mathcal{D}) and let $v: K(\mathcal{E}) \rightarrow \tilde{A}^*$ be a G -commuting labelling. Since every expansion of G is also critical, we may assume that G extends M

²The graph \mathcal{S} plays a crucial rôle in [10].

in the sense of Definition 2.4, i.e. $[x]_M = [xx^{-1}]_M$ for every word x for which $[x]_G = 1_G$. For every edge e of \mathcal{E} we define $u(e) := v(e)v(e)^{-1}w(\pi_e)$ (w and π_e as in Definition 3.6). If e is a loop edge then π_e is an empty path, and $[u(e)]_F = 1_F$ and $[u(e)]_M = [v(e)]_M$ since $[v(e)]_G = 1_G$ implies that $[v(e)]_M$ is idempotent. For non-loop edges (that is, edges of \mathcal{D}), $[v(e)]_M \leq [w(\pi_e)]_M$ implies $[v(e)]_M = [v(e)v(e)^{-1}w(\pi_e)]_M = [u(e)]_M$. For a closed path $e_1 \cdots e_t$ in \mathcal{E} then

$$[u(e_1 \cdots e_t)]_F = [w(\pi_{e_1} \cdots \pi_{e_t})]_F = 1_F,$$

since the path $\pi_{e_1} \cdots \pi_{e_t}$ is a closed path in the tree \mathcal{T} and therefore the corresponding reduced path is the empty path at $\alpha\pi_{e_1} = \alpha e_1$.

So far we have obtained Theorem 3.3 for finite connected graphs. Now let \mathcal{E} be any finite graph with connected components $\mathcal{E}_1, \dots, \mathcal{E}_n$. For every i there is a finite A -generated group G_i serving as in Theorem 3.3 for M and \mathcal{E}_i . Then the A -generated direct product $G := G_1 \times \cdots \times G_n$ serves as required for the inverse monoid M and the graph $\mathcal{E} = \mathcal{E}_1 \sqcup \cdots \sqcup \mathcal{E}_n$. \square

Remark 3.8. The above proof of Theorem 3.3 does not use the bound on the diameter of the tree \mathcal{T} in Definition 3.6 (critical group). That condition will be used, however, for the inductive procedure in the proof of Lemma 3.7.

Proof of Lemma 3.7 by use of Theorem 3.5. This proof is from [6, Section 7], but polished, with many details included and adjusted to the present context. In the first part we establish the claim for cycle graphs of length up to n in the rôle of \mathcal{E} , uniformly for all M . In the second part, this is lifted to the general case of connected graphs \mathcal{E} on n vertices, essentially by induction on the cyclomatic number: surgery that eliminates one simple cycle of \mathcal{E} reduces the claim to the claims for lower cyclomatic number, cycle graphs and trees. Altogether we derive Lemma 3.7 from Theorem 3.5, hence from Theorem 3.1.

To the first part. For any positive integer $n \geq 2$ let \mathcal{C}_n be the cycle graph of length n . From Theorem 3.5 we get the existence of a critical group for every pair (M, \mathcal{C}_n) . Indeed, let G be as in the statement of Theorem 3.5 and let $v: K(\mathcal{C}_n) \rightarrow \tilde{A}^*$ be a G -commuting labelling of \mathcal{C}_n . Denote the vertices of \mathcal{C}_n by o_1, \dots, o_n and consider edges $e_i: o_{i-1} \rightarrow o_i$ (indices to be taken mod n). Then $[v_1 \cdots v_n]_G = 1_G$ for $v_i = v(e_i)$. According to Theorem 3.5 there are words u_1, \dots, u_n such that $[u_i]_M = [v_i]_M$ and $[u_1 \cdots u_n]_F = 1_F$. The claim that G is critical for the pair (M, \mathcal{C}_n) follows from the following lemma. Note that considering labellings v of \mathcal{C}_n satisfying $v(e_{t+1}) = \cdots = v(e_n) = 1$ (the empty word) we find that G is critical also for every pair (M, \mathcal{C}_t) with $t < n$.

Lemma 3.9. *Let u_1, \dots, u_n be words such that $[u_1 \cdots u_n]_F = 1_F$ and let o_1, \dots, o_n be the vertices of the cycle graph \mathcal{C}_n of length n . Then there exists a finite tree \mathcal{T} of diameter at most $n - 1$ with $V(\mathcal{T}) \supseteq \{o_1, \dots, o_n\}$ and a labelling $w: K(\mathcal{T}) \rightarrow \tilde{A}^*$ such that $[u_i]_M \leq [w(\pi_i)]_M$ where π_i is the reduced path $o_{i-1} \rightarrow o_i$ in \mathcal{T} (indices to be taken mod n).*

Proof. Consider first the special case that the elements $[u_1 \cdots u_i]_F$ ($i = 1, \dots, n$) are pairwise distinct; this means that for every proper factor $u_i \cdots u_j$ of $u_1 \cdots u_n$ we have $[u_i \cdots u_j]_F \neq 1_F$. Let \mathcal{C}_n be the cycle graph of length n with edges $e_i: o_{i-1} \rightarrow o_i$ (indices to be taken mod n) and labelling $u: e_i \mapsto u_i$. For every i let $v_i := \text{red}(u_i)$. Note that no word v_i is empty. Now form a new graph \mathcal{R} by replacing in \mathcal{C}_n every edge $e_i: o_{i-1} \rightarrow o_i$ by an arc of the form

$$\vartheta_i: o_{i-1} \rightarrow p_{i1} \rightarrow \cdots \rightarrow p_{i,k_i-1} \rightarrow o_i \quad (3.4)$$

of length k_i the length of the word v_i , with new vertices p_{ij} and the edges being labelled by the letters of v_i so that $\ell(\vartheta_i) = v_i$. Assume that distinct arcs ϑ_i and ϑ_j have no edges and no intermediate vertices p_{ir} and p_{js} in common. The result is an A -labelled cycle graph of length $|v_1 \cdots v_n|$. Next apply Stallings foldings to get the maximum A -graph quotient $\mathcal{W} := \mathcal{R}/\Theta$ of \mathcal{R} . No two distinct vertices o_i and o_j of \mathcal{R} are identified by the folding process since in \mathcal{R} there is no path $o_i \rightarrow o_j$ labelled by a Dyck word. Hence $\{o_1, \dots, o_n\}$ can indeed be realised as a subset of $V(\mathcal{W})$. In the folding process no two vertices of any of the arcs (3.4) are identified with each other since no two such vertices are connected by a path labelled by a Dyck word. It follows that in \mathcal{W} we have a path $o_{i-1} \rightarrow o_i$ labelled v_i for every i . Since $v_1 \cdots v_n$ is a Dyck word, the folding process transforms the cycle graph \mathcal{R} into a tree: this follows easily by induction on the length of the word $v_1 \cdots v_n$. That is, \mathcal{W} is a finite tree in which the label of the reduced path $o_{i-1} \rightarrow o_i$ is $v_i = \text{red}(u_i)$, and clearly we have $[u_i]_M \leq [\text{red}(u_i)]_M$. We need to modify \mathcal{W} a bit in order to get the required bound on the diameter of \mathcal{T} . For this purpose, consider all arcs spanned by simple paths

$$q_0 \xrightarrow{f_1} q_1 \xrightarrow{f_2} \cdots \xrightarrow{f_{t-1}} q_{t-1} \xrightarrow{f_t} q_t \quad (3.5)$$

of consecutive edges f_1, \dots, f_t in \mathcal{W} such that the intermediate vertices q_1, \dots, q_{t-1} are all of degree 2 and none is of the form o_i while each of the vertices q_0 and q_t is of degree at least 3 or is of the form o_i for some i . If some reduced path between any two vertices o_i and o_j in \mathcal{W} traverses one of the edges f_s of (3.5) then it needs to traverse the entire arc (3.5). Hence we may contract the path (3.5) to a single edge, that is, replace every arc spanned by a path of the form (3.5) by the single edge $f: q_0 \rightarrow q_t$ (and its inverse) and label f accordingly, that is, $w(f)$ is the string of the (letter) labels of the edges f_1, \dots, f_t ; for all other edges f , let $w(f)$ be the label induced by the labelling of \mathcal{W} . In the resulting tree \mathcal{T} then, the label of every reduced path $o_{i-1} \rightarrow o_i$ is the same as in \mathcal{W} (namely $v_i = \text{red}(u_i)$). The tree \mathcal{T} now has at most n vertices of degree up to 2, hence has diameter at most $n - 1$ as required by Definition 3.6 (critical group).

The general case, namely that some of the elements $[u_1 \cdots u_i]_F$ ($i = 1, \dots, n$) coincide, can be done by induction on n where the case $n = 1$ is trivial. So let $n > 1$ and assume the claim of the lemma to be true for all products $v_1 \cdots v_k$ of length k smaller than n , and let u_1, \dots, u_n with

$[u_1 \cdots u_n]_F = 1_F$ be given. If $[u_i \cdots u_j]_F \neq 1_F$ for every proper factor $u_i \cdots u_j$ of $u_1 \cdots u_n$, then we are done, by the argument above. So let us assume that $[u_i \cdots u_j]_F = 1_F$ for some proper factor $u_i \cdots u_j$ of $u_1 \cdots u_n$. It is sufficient — possibly after some cyclic permutation of the words u_1, \dots, u_n and some change of indices — to consider the case $[u_1 \cdots u_k]_F = 1_F$ for some $k < n$. Then also $[u_{k+1} \cdots u_n]_F = 1_F$. By the inductive assumption, there exist finite trees \mathcal{T}_1 and \mathcal{T}_2 with $V(\mathcal{T}_1) \supseteq \{o_1, \dots, o_k\}$ and $V(\mathcal{T}_2) \supseteq \{o_{k+1}, \dots, o_n\}$, of diameters at most $k-1$ and $n-k-1$, respectively, and labellings $w_i: K(\mathcal{T}_i) \rightarrow \tilde{A}^*$ ($i = 1, 2$) according to the claim. This means that (i) for $1 \leq i \leq k$, if ϱ_i is the reduced path $o_{i-1} \rightarrow o_i$ in \mathcal{T}_1 (indices to be taken mod k), then $[u_i]_M \leq [w_1(\varrho_i)]_M$ and (ii) for $k+1 \leq i \leq n$, if ς_i is the reduced path $o_{i-1} \rightarrow o_i$ in \mathcal{T}_2 (indices to be taken mod $n-k$), then $[u_i]_M \leq [w_2(\varsigma_i)]_M$. We define the tree \mathcal{T} by

$$\mathcal{T} := \mathcal{T}_1 \sqcup \mathcal{T}_2 \sqcup \{f, f^{-1}\}$$

with $f: o_k \rightarrow o_n$ a new edge connecting the vertices o_k of \mathcal{T}_1 and o_n of \mathcal{T}_2 . We set $w_f(f^{\pm 1}) := 1$ and

$$w := w_1 \sqcup w_2 \sqcup w_f$$

to get a labelling of \mathcal{T} . We note that in case (i) above the reduced path $o_{i-1} \rightarrow o_i$ in \mathcal{T} coincides with the one in \mathcal{T}_1 provided that $i > 1$ while in case (ii) the reduced path $o_{i-1} \rightarrow o_i$ in \mathcal{T} coincides with the one in \mathcal{T}_2 provided that $i > k+1$. This implies that $[u_i]_M \leq [w(\pi_i)]_M$ for $i = 2, \dots, k, k+2, \dots, n$ where π_i is the reduced path $o_{i-1} \rightarrow o_i$ in \mathcal{T} . The reduced path $\pi_1: o_n \rightarrow o_1$ in \mathcal{T} is given by $\pi_1 = f^{-1}\varrho_1$ while the reduced path $\pi_{k+1}: o_k \rightarrow o_{k+1}$ in \mathcal{T} is $\pi_{k+1} = f\varsigma_{k+1}$. The inequality $[u_i]_M \leq [w(\pi_i)]_M$ therefore holds also for $i = 1$ and $i = k+1$, as required. Since, by the inductive assumption, \mathcal{T}_1 has diameter at most $k-1$ and \mathcal{T}_2 has diameter at most $n-k-1$, \mathcal{T} has diameter at most $n-1$. \square

To the second part. We consider all connected graphs \mathcal{E} without loop edges on n vertices. We show the existence of a critical group for every pair (M, \mathcal{E}) by induction on the number l of geometric edges of \mathcal{E} , simultaneously for all M . (In effect, the induction is on the cyclomatic number.) The base case of the induction is $l = n-1$, in which case the graph itself is a tree; we may take $\mathcal{T} = \mathcal{E}$, $w = v$ and G any group, in particular any group G that extends M in the sense of Definition 2.4. So let $l > n-1$ and suppose that the claim is true for all connected graphs \mathcal{D} on n vertices with fewer than l geometric edges and all finite A -generated inverse monoids N . Let \mathcal{E} be given with l geometric edges, let k be an edge of \mathcal{E} that is contained in a non-trivial cycle subgraph of \mathcal{E} , and let $\mathcal{D} = \mathcal{E} \setminus \{k, k^{-1}\}$. Note that $V(\mathcal{E}) = V(\mathcal{D})$. With the first part of the proof we obtain a group G critical for (M, \mathcal{C}_n) and hence, as remarked at the end of the last paragraph before Lemma 3.9, also for (M, \mathcal{C}_t) for all $t \leq n$. By the inductive assumption, the claim of the lemma is true for any pair (N, \mathcal{D}) . So let H be an expansion of

G that is critical for the pair $(M \times G, \mathcal{D})$. We claim that H is critical for (M, \mathcal{E}) ; for later reference we state:

$$\begin{aligned} G \text{ critical for } (M, \mathcal{C}_n) \text{ and } H \text{ critical for } (M \times G, \mathcal{D}) \\ \implies H \text{ critical for } (M, \mathcal{E}). \end{aligned} \quad (3.6)$$

Let $v: K(\mathcal{E}) \rightarrow \tilde{A}^*$ be an H -commuting labelling of \mathcal{E} . We choose edges k_1, \dots, k_t in $\mathcal{D} = \mathcal{E} \setminus \{k, k^{-1}\}$ so that $k_1 \cdots k_t: \omega k \rightarrow \alpha k$ is a simple path (that is, there is no repetition of vertices) in \mathcal{D} and thus $kk_1 \cdots k_t$ is a simple closed path in \mathcal{E} . Since v is H -commuting,

$$[v(k)v(k_1 \cdots k_t)]_H = 1_H. \quad (3.7)$$

Since the labelling $v \upharpoonright \mathcal{D}$ is also H -commuting, there exists, by the inductive assumption, a finite tree \mathcal{X} of diameter at most $n - 1$ such that

$$V(\mathcal{D}) \subseteq V(\mathcal{X}) \quad (3.8)$$

and a labelling $w_{\mathcal{X}}: K(\mathcal{X}) \rightarrow \tilde{A}^*$ such that for every edge e of \mathcal{D} , if π_e is the reduced path $\alpha e \rightarrow \omega e$ in \mathcal{X} , then $[v(e)]_{M \times G} \leq [w_{\mathcal{X}}(\pi_e)]_{M \times G}$. This holds in particular for the edges k_i : $[v(k_i)]_{M \times G} \leq [w_{\mathcal{X}}(\pi_{k_i})]_{M \times G}$ for all i . As a consequence we get

$$[v(k_i)]_G = [w_{\mathcal{X}}(\pi_{k_i})]_G \text{ for } i = 1, \dots, t, \quad (3.9)$$

since the relation $[v]_{M \times G} \leq [w]_{M \times G}$ implies $[v]_G \leq [w]_G$ (G is a quotient of $M \times G$) and \leq is the identity relation in the group G .

Next, we look more carefully at the reduced path $\omega k \rightarrow \alpha k$ in \mathcal{X} . We write it explicitly as

$$\alpha k_1 = p_0 \xrightarrow{f_1} p_1 \xrightarrow{f_2} p_2 \cdots p_{s-1} \xrightarrow{f_s} p_s = \omega k_t$$

and record that $s \leq n - 1$ by the assumption on the diameter of \mathcal{X} . Let \mathcal{C} be the cycle graph spanned by the path $f_1 \cdots f_s$ augmented by an extra edge $f_0: p_s = \alpha k \rightarrow \omega k = p_0$ (and its inverse), so that \mathcal{C} is a cycle graph of length $s + 1 \leq n$. We endow \mathcal{C} with the labelling $w_{\mathcal{C}}$ defined by

$$w_{\mathcal{C}}(f_i) = \begin{cases} w_{\mathcal{X}}(f_i) & \text{for } i = 1, \dots, s \\ v(k) & \text{for } i = 0. \end{cases}$$

The path $f_1 \cdots f_s$ is the reduced form of the path $\pi_{k_1} \cdots \pi_{k_t}$, and hence

$$[w_{\mathcal{X}}(f_1 \cdots f_s)]_F = [w_{\mathcal{X}}(\pi_{k_1} \cdots \pi_{k_t})]_F$$

so that

$$[w_{\mathcal{X}}(f_1 \cdots f_s)]_G = [w_{\mathcal{X}}(\pi_{k_1} \cdots \pi_{k_t})]_G. \quad (3.10)$$

Furthermore, (3.7) implies $[v(k)v(k_1 \cdots k_t)]_G = 1_G$ since $H \twoheadrightarrow G$. In combination with (3.9) and (3.10) this implies that the labelling $w_{\mathcal{C}}$ is G -commuting. It follows that there exists a finite tree \mathcal{Y} such that $V(\mathcal{C}) \subseteq$

$V(\mathcal{Y})$, and a labelling $w_{\mathcal{Y}}: K(\mathcal{Y}) \rightarrow \tilde{A}^*$ such that, if ρ_i is the reduced path $p_{i-1} \rightarrow p_i$ in \mathcal{Y} , then

$$[w_{\mathcal{C}}(f_i)]_M \leq [w_{\mathcal{Y}}(\rho_i)]_M \text{ for } i = 0, \dots, s \quad (3.11)$$

(indices to be taken mod $s + 1$). We now remove the edges $f_i^{\pm 1}$ from the tree \mathcal{X} and form the disjoint union with \mathcal{Y} , the result of which is the forest

$$\mathcal{U} := \mathcal{Y} \sqcup \mathcal{X} \setminus \{f_1^{\pm 1}, \dots, f_s^{\pm 1}\}$$

consisting of the connected components \mathcal{Y} and $\mathcal{X}_0, \dots, \mathcal{X}_s$, where \mathcal{X}_i is the component of p_i in the forest $\mathcal{X} \setminus \{f_1^{\pm 1}, \dots, f_s^{\pm 1}\}$. Let $\mathcal{Z} := \mathcal{U}/\Theta$ where Θ is the congruence on \mathcal{U} that identifies every vertex p_i of \mathcal{X}_i with the vertex p_i in \mathcal{Y} . Restricted to the tree \mathcal{Y} and each of the trees \mathcal{X}_i , Θ is the identity relation. Then \mathcal{Z} is a tree and we may consider \mathcal{Y} and every \mathcal{X}_i as subgraphs of \mathcal{Z} and such that $\mathcal{Y} \cap \mathcal{X}_i = \{p_i\}$, and we identify each vertex with its Θ -class. In particular, we have the inclusions

$$V(\mathcal{E}) = V(\mathcal{D}) \subseteq V(\mathcal{X}) \subseteq V(\mathcal{Z}).$$

Figure 2 gives an overview of the various graphs involved.




FIGURE 2. The trees \mathcal{X} , \mathcal{Y} and \mathcal{Z} and the cycle graph \mathcal{C} (dotted)

On the set of edges, Θ is the identity relation, whence we may write $K(\mathcal{Z}) = K(\mathcal{Y}) \cup \bigcup_{i=0}^s K(\mathcal{X}_i)$ (identifying each edge of \mathcal{U} with its Θ -class).

Next we endow \mathcal{Z} with its natural labelling: for any $f \in K(\mathcal{Z})$ let

$$w_{\mathcal{Z}}(f) = \begin{cases} w_{\mathcal{X}}(f) & \text{if } f \in \mathcal{X}_i \text{ for some } i \\ w_{\mathcal{Y}}(f) & \text{if } f \in \mathcal{Y}. \end{cases}$$

We want to show that the tree \mathcal{Z} with the labelling $w_{\mathcal{Z}}$ satisfies the requirements of Definition 3.6 for \mathcal{E} (with the exception of the condition on the diameter of the required tree). For this we show that, for every edge $e \in K(\mathcal{E})$ and the associated reduced path $\varsigma_e: \alpha e \rightarrow \omega e$ in \mathcal{Z} , we have $[v(e)]_M \leq [w_{\mathcal{Z}}(\varsigma_e)]_M$. We consider two cases: $e = k$ and $e \neq k^{\pm 1}$, that is, $e \in \mathcal{D}$.

Case 1: $e = k$. In this case, $\alpha k = p_s$ and $\omega k = p_0$ and the reduced path $\varsigma_k: \alpha k \rightarrow \omega k$ in \mathcal{Z} coincides with the one in \mathcal{Y} , which was earlier denoted ρ_0 . Since $v(k) = w_{\mathcal{C}}(f_0)$ we have

$$[v(k)]_M = [w_{\mathcal{C}}(f_0)]_M \leq [w_{\mathcal{Y}}(\rho_0)]_M = [w_{\mathcal{Z}}(\varsigma_k)]_M$$

by (3.11).

Case 2: $e \neq k^{\pm 1}$. First recall that

$$[v(e)]_{M \times G} \leq [w_{\mathcal{X}}(\pi_e)]_{M \times G} \quad (3.12)$$

by the construction of \mathcal{X} where π_e is the reduced path $\alpha e \rightarrow \omega e$ in \mathcal{X} . If π_e does not use any of the edges $f_1^{\pm 1}, \dots, f_s^{\pm 1}$, then the reduced path $\varsigma_e: \alpha e \rightarrow \omega e$ in \mathcal{Z} coincides with π_e and we are done. But it may happen that the path π_e traverses a segment of the arc spanned by $f_1 \cdots f_s$, that is,

$$\pi_e = \pi_1 f_r \cdots f_t \pi_2 \quad \text{or} \quad \pi_e = \pi_2 f_t^{-1} \cdots f_r^{-1} \pi_1 \quad (3.13)$$

for some $r \leq t$ and some reduced paths π_1 and π_2 in \mathcal{X} . It follows that π_1 runs in \mathcal{X}_{r-1} and π_2 runs in \mathcal{X}_t (recall that $f_i: p_{i-1} \rightarrow p_i$); it is sufficient to deal with the first case in (3.13). The reduced path $\varsigma_e: \alpha e \rightarrow \omega e$ in \mathcal{Z} then is given by $\varsigma_e = \pi_1 \rho \pi_2$ where ρ is the reduced path $p_{r-1} \rightarrow p_t$ in \mathcal{Y} . Recall that the reduced path $p_{i-1} \rightarrow p_i$ in \mathcal{Y} is denoted ρ_i . Then the path ρ is the reduced form of the product $\rho_r \cdots \rho_t$, so that

$$[w_{\mathcal{Y}}(\rho_r \cdots \rho_t)]_I \leq [w_{\mathcal{Y}}(\rho)]_I \quad (3.14)$$

in every inverse monoid I . For any $r \leq i \leq t$ we have $w_{\mathcal{X}}(f_i) = w_{\mathcal{C}}(f_i)$ and by construction of \mathcal{Y} , $[w_{\mathcal{C}}(f_i)]_M \leq [w_{\mathcal{Y}}(\rho_i)]_M$, hence $[w_{\mathcal{X}}(f_i)]_M \leq [w_{\mathcal{Y}}(\rho_i)]_M$. Using this and (3.14) we get

$$[w_{\mathcal{X}}(f_r \cdots f_t)]_M \leq [w_{\mathcal{Y}}(\rho_r \cdots \rho_t)]_M \leq [w_{\mathcal{Y}}(\rho)]_M = [w_{\mathcal{Z}}(\rho)]_M.$$

Since π_1 and π_2 are reduced paths in \mathcal{X} , $w_{\mathcal{X}}(\pi_i) = w_{\mathcal{Z}}(\pi_i)$ for $i = 1, 2$. Putting it all together we get

$$\begin{aligned} [w_{\mathcal{Z}}(\varsigma_e)]_M &= [w_{\mathcal{Z}}(\pi_1) \cdot w_{\mathcal{Z}}(\rho) \cdot w_{\mathcal{Z}}(\pi_2)]_M = [w_{\mathcal{X}}(\pi_1) \cdot w_{\mathcal{Y}}(\rho) \cdot w_{\mathcal{X}}(\pi_2)]_M \\ &= [w_{\mathcal{X}}(\pi_1)]_M [w_{\mathcal{Y}}(\rho)]_M [w_{\mathcal{X}}(\pi_2)]_M \\ &\geq [w_{\mathcal{X}}(\pi_1)]_M [w_{\mathcal{X}}(f_r \cdots f_t)]_M [w_{\mathcal{X}}(\pi_2)]_M \\ &= [w_{\mathcal{X}}(\pi_1 f_r \cdots f_t \pi_2)]_M = [w_{\mathcal{X}}(\pi_e)]_M \geq [v(e)]_M, \end{aligned}$$

where the last inequality follows from (3.12) since $M \times G \twoheadrightarrow M$.

For H to be critical for the pair (M, \mathcal{E}) one little detail is still missing: we need to modify \mathcal{Z} so that the diameter of the resulting tree is at most $n - 1$. To this end we proceed just as we did at the beginning of the proof: consider in \mathcal{Z} all maximal arcs spanned by simple paths of the form (3.5) with intermediate vertices of degree 2 but none of it being some $p \in V(\mathcal{E})$; replace every such arc by a single edge $e: q_0 \rightarrow q_t$ and define the label of each such edge by $w_{\mathcal{T}}(e) := w_{\mathcal{Z}}(e_1) \cdots w_{\mathcal{Z}}(e_t)$; for every edge e of \mathcal{Z} not involved in such an arc we just set $w_{\mathcal{T}}(e) := w_{\mathcal{Z}}(e)$. The resulting tree \mathcal{T} with the labelling $w_{\mathcal{T}}$ then fulfils all conditions required for Definition 3.6 (critical group). \square

3.3. Ash implies Herwig–Lascar. Let σ be a finite relational signature, and let us call an element y of a σ -structure \mathfrak{Y} *relation-free* if it is not contained in any relational tuple (y_1, \dots, y_r) (for any $r \geq 1$) of \mathfrak{Y} and let $\text{rf}(\mathfrak{Y})$ denote the set of all relation-free elements of \mathfrak{Y} . We define the *weight* $w(\mathfrak{Y})$ of \mathfrak{Y} as the number of relation-free elements of \mathfrak{Y} plus the number of the relational tuples of \mathfrak{Y} ; that is, $w(\mathfrak{Y}) := |\text{rf}(\mathfrak{Y})| + \sum_{R \in \sigma} |R^{\mathfrak{Y}}|$. Let $\mathbf{C} := \mathbf{Excl}(\Sigma)$ be the class of σ -structures defined by forbidden homomorphisms from a finite collection Σ of finite σ -structures $\mathfrak{U} \in \Sigma$. In the proof given below that \mathbf{C} is an EPPA-class, the weight of these forbidden structures \mathfrak{U} will give a bound on the number of vertices of the graphs used in the application of Theorem 3.3. Before we are ready for the proof we still have to discuss how to find a group suitable for the situation of Lemma 2.6.

Lemma 3.10. *Let M be a finite A -generated inverse monoid, $r \geq 1$ and let \mathcal{E} be the graph consisting of two vertices $o \neq p$ linked by r geometric edges. Let G be a finite A -generated group according to Theorem 3.3 with respect to M and \mathcal{E} . Then for every r -tuple w_1, \dots, w_r of words for which $[w_1]_G = \dots = [w_r]_G$ there exists a word w such that $[w_1]_M, \dots, [w_r]_M \leq [w]_M$.*

Proof. Let $k_1, \dots, k_r: o \rightarrow p$ be the edges from o to p . Label these edges according to $\ell(k_i^{\pm 1}) = w_i^{\pm 1}$ for all i . The assumption that $[w_1]_G = \dots = [w_r]_G$ implies that this labelling is G -commuting. By Theorem 3.3, there exist words v_i ($i = 1, \dots, r$) such that (i) $[v_i v_j^{-1}]_F = 1_F$ and (ii) $[w_i]_M = [v_i]_M$ for all i, j . From (i) it follows that $[v_i]_F = [v_j]_F$ for all i, j so that all words v_i have the same reduced form: $\text{red}(v_i) = \text{red}(v_j) =: w$. Hence $[w]_M \geq [v_i]_M = [w_i]_M$ for all i . \square

We are now ready to derive the Theorem of Herwig–Lascar from Ash’s Theorem.

Proof of Theorem 3.2 by use of Theorem 3.3. Let \mathbf{C} be a class of σ -structures defined by finitely many forbidden homomorphisms of finite σ -structures, let $\mathfrak{X} = (X; (R^{\mathfrak{X}})_{R \in \sigma})$ be a finite member of \mathbf{C} equipped with some set \tilde{A} of partial automorphisms, and let \mathcal{P} be the A -graph with vertex set X encoding these partial mappings. For the set Σ of forbidden structures \mathfrak{U} let n be

the maximum of the weights of all members of Σ , that is, $n = \max_{\mathfrak{U} \in \Sigma} w(\mathfrak{U})$. Now choose a finite A -generated group G according to Theorem 3.3 for the inverse monoid $M = \mathcal{T}(\mathcal{P})$ and for the graph \mathcal{P} consisting of the disjoint union of all (isomorphism types of) connected graphs having at most n vertices and, between any two distinct vertices, at most $|X|$ many geometric edges. By Lemma 3.10 in combination with Lemma 2.6, \mathfrak{X} is an induced substructure of $\mathfrak{X}_{\mathcal{P}}G$. We shall prove that, if $\mathfrak{X}_{\mathcal{P}}G$ is not \mathfrak{U} -free for some $\mathfrak{U} \in \Sigma$, then neither is the free extension $\mathfrak{X}_{\mathcal{P}}F$.

Let $\mathfrak{U} \in \Sigma$ and let $\phi: \mathfrak{U} \rightarrow \mathfrak{X}_{\mathcal{P}}G$ be a homomorphism. The weight of $\phi(\mathfrak{U})$ as a weak substructure of $\mathfrak{X}_{\mathcal{P}}G$ is at most n . So there exists a subset $J \subseteq G$ of size n such that every relational tuple $(\phi(t_1), \dots, \phi(t_r))$ of $\phi(\mathfrak{U})$ is completely contained in some single translate Xg for some $g \in J$ (meaning that every entry $\phi(t_i)$ of this tuple is contained in that translate) and that $\phi(\mathfrak{U}) \subseteq \bigcup_{g \in J} Xg$. In other words, by the definition of the weight, the size n of J guarantees that (i) the members of every relational tuple of $\phi(\mathfrak{U})$ are contained in at least one single translate Xg of X (one translate for each tuple) and (ii) the entire homomorphic image $\phi(\mathfrak{U})$ is contained in the union $\bigcup_{g \in J} Xg$.³ For convenience, we consider X as a subset of $X_{\mathcal{P}}G$ as well as of $X_{\mathcal{P}}F$, in fact we assume that $X = X_{\mathcal{P}}G \cap X_{\mathcal{P}}F$. We want to establish a homomorphism $\phi(\mathfrak{U}) \rightarrow \mathfrak{X}_{\mathcal{P}}F$. The idea is to find a function $\psi: \bigcup_{g \in J} Xg \rightarrow X_{\mathcal{P}}F$ which is a homomorphism of σ -structures when restricted to $\phi(\mathfrak{U})$. More precisely, for each $g \in J$ we would like to choose some $f_g \in F$ and consider the mapping:

$$\begin{array}{ccccc} Xg & \longrightarrow & X & \longrightarrow & Xf_g \subseteq X_{\mathcal{P}}F \\ z & \longmapsto & zg^{-1} & \longmapsto & (zg^{-1})f_g \end{array} \quad (3.15)$$

where g^{-1} acts on $X_{\mathcal{P}}G$ and maps Xg to X , and f_g acts on $X_{\mathcal{P}}F$ and maps X to Xf_g . Suppose we had a mapping $\psi: \bigcup_{g \in J} Xg \rightarrow X_{\mathcal{P}}F$ whose restriction $\psi \upharpoonright Xg$ to any translate Xg can be described by (3.15). Such a mapping would be compatible with the relational structure of $\phi(\mathfrak{U})$, hence a homomorphism when restricted to $\phi(\mathfrak{U})$. Indeed, any relational tuple $(\phi(t_1), \dots, \phi(t_r)) \in R^{\mathfrak{X}_{\mathcal{P}}G}$ (for some $R \in \sigma$) of $\phi(\mathfrak{U})$ is contained in some translate Xg . The mapping described in (3.15) certainly maps the elements $\phi(t_1), \dots, \phi(t_r)$ to elements of Xf_g , and the image tuple under ψ belongs to $R^{\mathfrak{X}_{\mathcal{P}}F}$ — this follows from the definition of the σ -structures $\mathfrak{X}_{\mathcal{P}}G$ and $\mathfrak{X}_{\mathcal{P}}F$.

The difficulty lies in the well-definedness of the map ψ according to (3.15). Distinct translates Xg and Xh for $g, h \in J$ may overlap: the choices $f_g, f_h \in F$ should satisfy

$$(zg^{-1})f_g = (zh^{-1})f_h$$

³At this point we would not require to consider relation-free elements, since we could redefine the homomorphism ϕ to map the relation-free elements wherever we wish. The present pedantic approach is in a sense rather a preparation for the proof of Theorem 4.18.

for all $g, h \in J$ for which $Xg \cap Xh \neq \emptyset$ and all $z \in Xg \cap Xh$. This would make such a mapping ψ well-defined. The situation is depicted in Figure 3.




FIGURE 3

This is the point where we apply Ash's Theorem. Define a graph \mathcal{J} with set of vertices $V(\mathcal{J}) = J$ in which two vertices g and h are connected by one or more edges if and only if $Xg \cap Xh \neq \emptyset$. In fact, there will be $|Xg \cap Xh|$ many edges from g to h , and their inverses. Suppose that $Xg \cap Xh = \{z_1, \dots, z_r\}$; then there are $x_1, \dots, x_r, y_1, \dots, y_r \in X$ such that $x_i g = z_i = y_i h$ for all i , hence $x_i g h^{-1} = y_i$ for all i . According to Lemma 2.2 there exist words $v_1, \dots, v_r \in \tilde{A}^*$ such that $x_i g h^{-1} = y_i = x_i \cdot [v_i]_M$ for all i and $[v_i]_G = g h^{-1}$. In this situation we impose edges $e_1, \dots, e_r: g \rightarrow h$ with labels v_1, \dots, v_r , respectively. The corresponding inverse edges $e_i^{-1}: h \rightarrow g$ then have labels v_i^{-1} and satisfy $y_i h g^{-1} = x_i = y_i \cdot [v_i^{-1}]_M$. Doing this for all pairs $g, h \in J$ for which $Xg \cap Xh \neq \emptyset$ yields a word labelling of \mathcal{J} that commutes over G : since for every edge $e: g \rightarrow h$, the G -value of its label is gh^{-1} . Hence the label of any closed path has G -value 1_G . According to Theorem 3.3 there exists an M -related relabelling that commutes over the free group F . For the edges $e_1, \dots, e_r: g \rightarrow h$ and their labels v_1, \dots, v_r (mentioned above) this means that there are words $w_1, \dots, w_r \in \tilde{A}^*$ such that $[v_i]_M = [w_i]_M$ and the overall labelling by these words commutes over F . In particular, the words w_i (of the edges $e_i: g \rightarrow h$) all have the same reduced form, say w , and $[v_i]_M = [w_i]_M \leq [w]_M$ so that also $x_i \cdot [w]_M = y_i$ for all i . We relabel the graph once more by replacing every w_i by this reduced word $w =: w^{(g,h)}$. Again this is done for all pairs g, h for which $Xg \cap Xh \neq \emptyset$.

Suppose first that \mathcal{J} is connected. Choose a base vertex o and define, for any vertex g , an element f_g by setting $f_g := [p]_F$ where p is the label of some path $g \rightarrow o$; this labelling is meant with respect to the last relabelling where every edge $e: g \rightarrow h$ is labelled $w^{(g,h)}$. Since the labelling commutes over F , the element f_g is well-defined and does not depend on the choice of the path $g \rightarrow o$. Now let $z \in Xg \cap Xh$ for some $g, h \in J$. Then there are $x, y \in X$ such that $xg = z = yh$, hence $x = zg^{-1}$, $y = zh^{-1}$ and $xgh^{-1} = y = x \cdot [w^{(g,h)}]_M$. From the definition of f_- it follows that $f_g = [w^{(g,h)}]_F f_h$. Altogether we arrive at

$$(zg^{-1})f_g = xf_g = x[w^{(g,h)}]_F f_h = (x \cdot [w^{(g,h)}]_M)f_h = yf_h = (zh^{-1})f_h,$$

which shows that the mapping $\psi: \bigcup_{g \in J} Xg \rightarrow X_{\mathcal{P}}F$ according to (3.15) is well-defined. If the graph \mathcal{J} is not connected, then the same argument can be applied to each connected component of \mathcal{J} separately. It follows that the free extension $\mathfrak{X}_{\mathcal{P}}F$ is not \mathfrak{U} -free, and the theorem is proved. \square

3.4. Herwig–Lascar implies Ash. In this subsection we give a direct proof of the implication Herwig–Lascar \implies Ash, which seems to be interesting in its own right. Essentially it is a generalisation of the arguments in Section 3.1 from cycle graphs to arbitrary graphs; the proof of Proposition 3.4 can be seen as a warm up to the proofs of Lemmas 3.12 and 3.13. It also provides an essential ingredient for the understanding of the discussion of the group-theoretic version(s) of the Herwig–Lascar Theorem in Section 3.5. Otherwise this section is independent of the remainder of the paper.

Let $\mathcal{E} = (V \cup K; \alpha, \omega, {}^{-1})$ be a graph; a *coset labelling* of \mathcal{E} is a mapping $\ell: K \rightarrow 2^F$, $k \mapsto \ell(k) := H_k g_k$, with H_k a finitely generated subgroup of F and $g_k \in F$ such that $\ell(k^{-1}) = \ell(k)^{-1}$, that is, $H_{k^{-1}} g_{k^{-1}} = (H_k g_k)^{-1} = (g_k^{-1} H_k) g_k^{-1}$, which is the set of all inverses of the elements of $H_k g_k$. A coset labelling ℓ of \mathcal{E} *commutes over the A -generated group G* (is *G -commuting* for short) if there exists a G -commuting word labelling l such that $l(k) \in \ell(k)$ for every $k \in K$ (in this subsection we consider F to be the set of all reduced words in \tilde{A}^* and assume that every word labelling takes values in F). In a sense, a coset labelling is G -commuting if it *covers* a G -commuting word labelling. The essential achievement in this subsection will be to prove that Proposition 3.11 below is a consequence of the Herwig–Lascar Theorem.

Proposition 3.11. *Let \mathcal{E} be a finite graph and let ℓ be a coset labelling of \mathcal{E} . If ℓ is not F -commuting then ℓ is not G -commuting for some finite A -generated group G .*

Intuitively, the statement means that the non-commutativity with respect to F of a coset labelling of a finite graph can be witnessed (or recognised) by some finite group G ; or, to say it the other way round: the free group is well approximable by finite groups with respect to commutativity of coset labellings of finite graphs.

We first note that Proposition 3.11 is a direct generalisation of Proposition 3.4 from cycle graphs to arbitrary graphs. Indeed, let H_1, \dots, H_n be finitely generated subgroups of F and $g_1, \dots, g_n \in F$. The sequence of cosets H_1g_1, \dots, H_ng_n can be seen as a coset labelling of the cycle graph of length n . To say that $1 \notin H_1g_1 \cdots H_ng_n$ is equivalent to saying that this labelling is not F -commuting; to say that $1 \notin NH_1g_1 \cdots NH_ng_n$ is equivalent to saying that the original coset labelling is not G -commuting for the group $G = F/N$.

Next we note that it suffices to prove Proposition 3.11 (i.e. derive it from Theorem 3.2) for finite connected graphs without loop edges. Let \mathcal{E} be a finite connected graph and let \mathcal{D} be the graph obtained from \mathcal{E} by removing all loop edges (if any). Consider a coset labelling $\ell_{\mathcal{E}}$ that is not F -commuting. If the restriction of $\ell_{\mathcal{E}}$ to \mathcal{D} is also not F -commuting we are done. Suppose we have a labelling $\ell_{\mathcal{E}}$ of \mathcal{E} that is not F -commuting while its restriction to \mathcal{D} is F -commuting. Then there exists a loop edge e whose label $\ell_{\mathcal{E}}(e)$ is a proper coset $Hg \neq H$ (that is, $1_F \notin Hg$). There is a finite A -generated group G for which $[g]_G \neq 1_G$: G can be chosen to be the transition group of any finite completion $\overline{\mathcal{S}_{Hg}}$ of the Stallings graph \mathcal{S}_{Hg} (this is Stallings' proof [35] of Hall's Theorem [18]). Hence, given a labelling $\ell_{\mathcal{E}}$ of \mathcal{E} , if there exists a finite group G detecting that its restriction to \mathcal{D} is not F -commuting, then there also exists a finite group H detecting that $\ell_{\mathcal{E}}$ itself is not F -commuting. Finally, a graph \mathcal{E} satisfies the condition formulated in Proposition 3.11 if and only every connected component of \mathcal{E} does.

The intuitive idea of the proof roughly is as follows. Suppose we are given a graph \mathcal{E} and a coset labelling ℓ of \mathcal{E} . Let $e = \{k, k^{-1}\}$ be a geometric edge with incident vertices $u = \alpha k$ and $v = \omega k$. The idea is to replace e with the plain Schreier graph $\mathcal{S}_e := \mathcal{S}_{H_k g_k}^{\circ}$ where $H_k g_k = \ell(k)$, and to encode the entire labelling ℓ on the disjoint union $\mathcal{S} := \bigsqcup \mathcal{S}_e$ where the union runs over all geometric edges e of \mathcal{E} . In order that the latter union encodes the labelling completely we need to mark the two distinguished vertices of $\mathcal{S}_{H_k g_k}$ within \mathcal{S}_e so as to associate the correct coset $H_k g_k$ with the edge e when traversed in the direction from u to v . This can be achieved by *colouring* the initial vertex ι of $\mathcal{S}_{H_k g_k}$ by colour u and the terminal vertex τ of $\mathcal{S}_{H_k g_k}$ by colour v . When e is traversed in the opposite direction this colouring already provides the correct coset, namely $(g_k^{-1}H)g_k^{-1}$. Note that it may happen that a distinguished vertex carries both colours u and v (this happens exactly if the corresponding coset $H_k g_k$ coincides with the group H_k .) So far, on each connected component \mathcal{S}_e of \mathcal{S} , two vertices ι and τ of \mathcal{S}_e (possibly identical) are — thanks to the colouring — associated with the two vertices u and v of \mathcal{E} (namely those incident with the geometric edge e). From point of view of the vertices of \mathcal{E} we have, for every vertex u of \mathcal{E} , u -coloured vertices in \mathcal{S} , one in each component \mathcal{S}_e , for every geometric edge e incident with u . Now recall that the free group F acts on \mathcal{S} leaving invariant every component \mathcal{S}_e . For every vertex v of \mathcal{E} we consider the collection of all v -coloured distinguished vertices as being “rigidly entangled” forming a “ v -block” that

is moved “as a whole” by the action of F (although they belong to pairwise distinct components \mathcal{S}_e). Later, in the formal proof, these v -blocks will be realised as relational tuples whose entries are comprised of the v -coloured vertices (the sets of these will be denoted by M_v).

The idea now is that the original labelling ℓ commutes over F if and only if the action of F moves the v -blocks in such a way that, for every geometric edge e , the two distinguished vertices in \mathcal{S}_e are mapped to a common vertex of \mathcal{S}_e . (A trivial special case of this is when all labels $\ell(k)$ are subgroups: then all pairs of distinguished vertices in every component \mathcal{S}_e coincide from the very beginning and the labelling trivially commutes since every label contains the identity element 1_F .) More precisely, for every vertex v of \mathcal{E} there is an element $w_v \in F$ such that the following holds for the collection $(w_v : v \in V)$: for every geometric edge e of \mathcal{E} with incident vertices u and v the image under w_u of the u -coloured vertex of \mathcal{S}_e coincides with the image under w_v of the v -coloured vertex. It then follows that $w_u w_v^{-1}$ labels a path from the u -vertex to the v -vertex in \mathcal{S}_e . Consequently, the reduced form $\text{red}(w_u w_v^{-1})$ belongs to the coset label of the edge e (the label when the edge is traversed from u to v). It follows that, from any coset label we may select an element such that for any cycle in the graph \mathcal{E} , the product over these selected elements is the identity element 1_F . For the special case of a cycle graph \mathcal{E} this is the essence of what happens in the proof of Proposition 3.4. This idea will now be formalised and translated into the language of relational structures.

In the following we shall assume that the graph $\mathcal{E} = (V \cup K; \alpha, \omega, {}^{-1})$ in question is finite, connected and without loop edges. We introduce a relational signature $\sigma_V = \{R_v : v \in V\}$ consisting of relational symbols R_v , indexed by the vertices of \mathcal{E} , the arity of R_v being r_v , the degree of the vertex v . Let \mathfrak{C} and \mathfrak{C}^* be two σ_V -structures defined as follows. The base set of \mathfrak{C} is K , the set of edges of \mathcal{E} ; for $v \in V$ we let $K_v := \{k \in K : \omega k = v\}$. For every $v \in V$ we choose an enumeration k_1, \dots, k_{r_v} of K_v and set $R_v^{\mathfrak{C}} := \{(k_1, \dots, k_{r_v})\}$. This defines a σ_V -structure $\mathfrak{C} = (K; (R_v^{\mathfrak{C}})_{v \in V})$. We call \mathfrak{C} a σ_V -structure associated with the graph \mathcal{E} .

In the context of σ_V -structures, for any equivalence relation \sim on the base set K , the quotient structure \mathfrak{C}/\sim is defined on the quotient set K/\sim as follows: a tuple $(\bar{k}_1, \dots, \bar{k}_r)$ of the quotient set K/\sim is in $R_v^{\mathfrak{C}/\sim}$ if and only if for all i there are $q_i \in \bar{k}_i$ such that the tuple (q_1, \dots, q_r) is in $R_v^{\mathfrak{C}}$. Finally, given a σ_V -structure \mathfrak{C} associated with \mathcal{E} we define the σ_V -structure \mathfrak{C}^* by factoring the base set K of \mathfrak{C} by the equivalence relation \equiv which identifies every edge k with its inverse k^{-1} :

$$\mathfrak{C}^* := \mathfrak{C}/\equiv$$

where for $k_1, k_2 \in K$,

$$k_1 \equiv k_2 \iff k_1 = k_2 \text{ or } k_1 = k_2^{-1}.$$

We call \mathfrak{C}^* a *consolidated* σ_V -structure associated with the graph \mathcal{E} . In addition, any σ_V -structure \mathfrak{C}/\sim (including \mathfrak{C} itself) obtained by factoring the base set K of \mathfrak{C} by an equivalence relation \sim contained in \equiv (that is by identifying *some*, but not necessarily *all*, edges k with their inverses k^{-1}) is called an *approximation of \mathfrak{C}^** .

Prima facie, every structure \mathfrak{C} associated with a graph \mathcal{E} depends on the choice of the enumeration of each set K_v , but it is easy to see that any two different such structures \mathfrak{C}_1 and \mathfrak{C}_2 are isomorphic as σ_V -structures. However, this does not imply that the corresponding consolidated structures \mathfrak{C}_1^* and \mathfrak{C}_2^* are also isomorphic. The reason for this is that an isomorphism $\mathfrak{C}_1 \rightarrow \mathfrak{C}_2$ not necessarily respects the equivalence relation \equiv on K . An example \mathcal{E} for such a behaviour is the cycle graph on two vertices (that is, $|V| = 2$ and $|K| = 4$). In this case there exist four possible associated structures \mathfrak{C}_i , all of which are isomorphic as σ_V -structures, and also four possible consolidated structures \mathfrak{C}_i^* which belong to two different isomorphism classes of σ_V -structures. This peculiar behaviour is irrelevant for our purpose.

Recall the concepts of the Stallings graph \mathcal{S}_{Hg} and the Schreier graph \mathcal{S}_{Hg} of a coset Hg (see Section 2.6). Both are 2-pointed A -graphs consisting of the *plain* A -graph \mathcal{S}_{Hg}° (respectively \mathcal{S}_{Hg}°) (plain A -graph means that there are no distinguished vertices) and the distinguished (initial and terminal) vertices ι and τ . We use the suggestive notation

$$\mathcal{S}_{Hg} = (\iota, \mathcal{S}_{Hg}^\circ, \tau)$$

and set $\iota(\mathcal{S}_{Hg}) := \iota$ and $\tau(\mathcal{S}_{Hg}) := \tau$. For the “inverse” coset $(g^{-1}H)g^{-1}$ we have $\mathcal{S}_{Hg}^\circ = \mathcal{S}_{(g^{-1}H)g^{-1}}^\circ$ so that

$$\mathcal{S}_{(g^{-1}H)g^{-1}} = (\tau, \mathcal{S}_{Hg}^\circ, \iota),$$

which means that the plain A -graph \mathcal{S}_{Hg}° remains unchanged and just the two distinguished vertices ι and τ are swapped. This is expressed by the formula

$$\iota(\mathcal{S}_{Hg}) = \tau(\mathcal{S}_{(g^{-1}H)g^{-1}}) \text{ and } \iota(\mathcal{S}_{(g^{-1}H)g^{-1}}) = \tau(\mathcal{S}_{Hg}),$$

and the same holds for the corresponding Schreier graphs \mathcal{S}_{Hg} and $\mathcal{S}_{(g^{-1}H)g^{-1}}$.

Next consider a coset labelling ℓ of \mathcal{E} . For every edge k of \mathcal{E} we have $\ell(k) = H_k g_k$ for some finitely generated subgroup H_k of F and some $g_k \in F$ and such that $\ell(k^{-1}) = H_{k^{-1}} g_{k^{-1}} = (g_k^{-1} H_k) g_k^{-1}$, that is, $g_{k^{-1}} = g_k^{-1}$ and $H_{k^{-1}} = g_k^{-1} H_k$. Recall the notation for geometric edges $k^\bullet := \{k, k^{-1}\}$ and $K^\bullet := \{k^\bullet : k \in K\}$. For $e = k^\bullet$ let

$$\mathcal{S}_e := \mathcal{S}_{H_k g_k}^\circ = \mathcal{S}_{(g_k^{-1} H_k) g_k^{-1}}^\circ$$

be the plain Stallings graph of $H_k g_k$ — that is, the Stallings graph with the distinguished vertices ignored, which is also the plain Stallings graph of $H_{k^{-1}} g_{k^{-1}} = (g_k^{-1} H_k) g_k^{-1}$ — and let $\mathcal{S} := \bigsqcup_{e \in K^\bullet} \mathcal{S}_e$. We define a mapping $\varpi : K \rightarrow V(\mathcal{S})$ as follows: for $e = k^\bullet = \{k, k^{-1}\}$ let $\varpi(k) := \tau(\mathcal{S}_{H_k g_k})$ (the

terminal vertex of $\mathcal{S}_{H_k g_k} = (\iota, \mathcal{S}_e, \tau)$, considered as a vertex of $\mathcal{S}_e \subseteq \mathcal{S}$ and $\varpi(k^{-1}) := \iota(\mathcal{S}_{H_k g_k})$ (the initial vertex of $\mathcal{S}_{H_k g_k} = (\iota, \mathcal{S}_e, \tau)$, considered as a vertex of $\mathcal{S}_e \subseteq \mathcal{S}$). The discussion above about plain Stallings graphs implies that the mapping ϖ is well-defined in the sense that $\varpi(k^{-1}) = \tau(\mathcal{S}_{(g_k^{-1} H_k) g_k^{-1}}) = \tau(\mathcal{S}_{H_{k^{-1}} g_{k^{-1}}})$, the terminal vertex of the Stallings graph $\mathcal{S}_{H_{k^{-1}} g_{k^{-1}}}$.

If k_1 and k_2 are edges that belong to different geometric edges, that is, $k_1^\bullet \neq k_2^\bullet$, then $\varpi(k_1) \neq \varpi(k_2)$ since $\varpi(k_1) \in V(\mathcal{S}_{k_1^\bullet})$ while $\varpi(k_2) \in V(\mathcal{S}_{k_2^\bullet})$ (which are disjoint). In particular, the restriction $\varpi \upharpoonright K_v$ is injective for every $v \in V$. For later use we set $\varpi_v := \varpi \upharpoonright K_v$ and let $M_v := \varpi(K_v) = \varpi_v(K_v)$ and note that ϖ_v is a bijection $K_v \rightarrow M_v$ for every vertex v . Only in case when k_1 and k_2 belong to the same geometric edge $k_1^\bullet = k_2^\bullet$ the values $\varpi(k_1)$ and $\varpi(k_2)$ may coincide. This happens exactly if the initial and terminal vertices of $\mathcal{S}_{\ell(k_1)}$ (and hence also of $\mathcal{S}_{\ell(k_2)}$) coincide, which happens if and only if the coset $\ell(k_i)$ is actually a group.

Now we consider Schreier graphs; every Stallings graph \mathcal{S}_{H_g} is naturally embedded in the Schreier graph \mathcal{S}_{H_g} which, by definition, is the free completion $\mathcal{S}_{H_g} F$ of \mathcal{S}_{H_g} . For $e = \{k, k^{-1}\}$ we set $\mathcal{S}_e := \mathcal{S}_{H_k g_k}^\circ$ and $\mathcal{S} := \bigsqcup_{e \in K^\bullet} \mathcal{S}_e$. There are natural inclusions $\mathcal{S}_e \subseteq \mathcal{S}_e$ for every $e \in K^\bullet$ and hence also $\mathcal{S} \subseteq \mathcal{S}$. In particular, the mapping ϖ can be interpreted as a mapping $K \rightarrow V(\mathcal{S})$. Recall that the free group F acts by bijections on every set $V(\mathcal{S}_e)$ and hence on $V(\mathcal{S})$ on the right via $\rho \mapsto \rho w$, for every $\rho \in V(\mathcal{S})$ and every $w \in F$, and the set $V(\mathcal{S}_e)$ is invariant under this action for each $e \in K^\bullet$.

As in Section 3.1 we define a σ_V -structure on $V(\mathcal{S})$, that is, for every $v \in V$ we define a relation R_v of arity r_v on $V(\mathcal{S})$ as follows. First choose some σ_V -structure \mathfrak{C} associated with \mathcal{E} — in particular, for every $v \in V$, choose and fix an enumeration k_1, \dots, k_{r_v} for the elements of the set K_v ; then transfer this enumeration to the corresponding set M_v via the bijection ϖ_v :

$$\varpi_v k_1, \dots, \varpi_v k_{r_v}; \quad (3.16)$$

then set

$$(\varpi_v k_1, \dots, \varpi_v k_{r_v}) \in R_v \quad (3.17)$$

and add to this whatever is required so that the aforementioned action $\rho \mapsto \rho w$ by F becomes an automorphism of the σ_V -structure. More precisely

$$R_v = \{(\varpi_v k_1, \dots, \varpi_v k_{r_v})w : w \in F\}, \quad (3.18)$$

where the application of w to an r_v -tuple is understood componentwise. In this way we get a σ_V -structure $\mathfrak{S} = (S; (R_v)_{v \in V})$ on the base set $S = V(\mathcal{S})$ with relations R_v of arity r_v and the free group F acts on this structure by automorphisms on the right. The idea is to initialise a σ_V -structure on the subset $\bigcup_{v \in V} M_v$ (3.17) and then extend the structure to all of $V(\mathcal{S})$ via the action of F (3.18). The mapping $\varpi : K \rightarrow V(\mathcal{S})$ is a homomorphism of σ_V -structures $\mathfrak{C} \rightarrow \mathfrak{S}$. Its image $\varpi(K) = \bigcup_{v \in V} M_v$, as a (weak) substructure of \mathfrak{S} , forms an approximation of \mathfrak{C}^\star since for $k_1, k_2 \in K$, $\varpi(k_1) = \varpi(k_2)$

can happen only if $k_1 \equiv k_2$, as already mentioned. We state and prove the essential assertion about \mathfrak{S} .

Lemma 3.12. *If there is a homomorphism $\phi: \mathfrak{C}^* \rightarrow \mathfrak{S}$, then the coset labelling ℓ of \mathcal{E} commutes over F .*

Proof. In the following, the interpretations of the symbols R_v in \mathfrak{S} are written without superscript as in the previous text, while those in \mathfrak{C}^* are denoted as $R_v^{\mathfrak{C}^*}$. For every $v \in V$ we have $R_v = \{(\varpi_v k_1, \dots, \varpi_v k_{r_v})w : w \in F\}$. Let $\phi: \mathfrak{C}^* \rightarrow \mathfrak{S}$ be a homomorphism of σ_V -structures. Consider two vertices $u, v \in V$ connected by an edge in \mathcal{E} , say $k: u \rightarrow v$ and $k^{-1}: v \rightarrow u$. Let $K_v = \{k_1, \dots, k_{r_v}\}$ and $K_u = \{l_1, \dots, l_{r_u}\}$ with $\{(k_1^\bullet, \dots, k_{r_v}^\bullet)\} = R_v^{\mathfrak{C}^*}$ and $\{(l_1^\bullet, \dots, l_{r_u}^\bullet)\} = R_u^{\mathfrak{C}^*}$. By the definition of R_v and R_u , respectively, there exist $w_v, w_u \in F$ such that

$$\phi(k_1^\bullet, \dots, k_{r_v}^\bullet) = (\varpi_v k_1, \dots, \varpi_v k_{r_v})w_v$$

and

$$\phi(l_1^\bullet, \dots, l_{r_u}^\bullet) = (\varpi_u l_1, \dots, \varpi_u l_{r_u})w_u,$$

where the application of ϕ and $\cdot w_v$ respectively $\cdot w_u$ are understood componentwise. Note that the ordering of the entries in the respective tuples comes from (3.16). By construction, there are indices i, j such that $k = k_i$ and $k^{-1} = l_j$. But then $k_i^\bullet = l_j^\bullet$ which implies

$$(\varpi k)w_v = (\varpi_v k_i)w_v = \phi(k_i^\bullet) = \phi(l_j^\bullet) = (\varpi_u l_j)w_u = (\varpi k^{-1})w_u. \quad (3.19)$$

The equality 3.19 takes place in $V(\mathcal{S}_{k^\bullet})$ where \mathcal{S}_{k^\bullet} is the plain graph of the Schreiergraph $\mathcal{S}_{H_k g_k}$. By definition of the mapping ϖ we have $\varpi(k) = \tau := \tau(\mathcal{S}_{H_k g_k}) = \tau(\mathcal{S}_{H_k g_k})$ and $\varpi(k^{-1}) = \iota := \iota(\mathcal{S}_{H_k g_k}) = \iota(\mathcal{S}_{H_k g_k})$. Equality (3.19) then just says that $\iota w_u = \tau w_v$ so that $\tau = \iota w_u w_v^{-1}$. The latter equality holds in $\mathcal{S}_{H_k g_k}$, which implies that $\tau = \iota \cdot \text{red}(w_u w_v^{-1})$ holds in $\mathcal{S}_{H_k g_k}$, which in turn implies that $\text{red}(w_u w_v^{-1}) \in H_k g_k$. Observe that $u = \alpha k$ and $v = \omega k$. So we have shown that for every edge $k: \alpha k \rightarrow \omega k$, the word $w(k) := \text{red}(w_{\alpha k} w_{\omega k}^{-1})$ belongs to $\ell(k)$, and it is clear that the word labelling $k \mapsto w(k) = \text{red}(w_{\alpha k} w_{\omega k}^{-1})$ is F -commuting. \square

Next consider again, for $e \in K^\bullet$, the plain Stallings and Schreier graphs $\mathcal{S}_e \subseteq \mathcal{S}_e$ and their disjoint unions

$$\mathcal{S} := \bigsqcup_{e \in K^\bullet} \mathcal{S}_e \subseteq \bigsqcup_{e \in K^\bullet} \mathcal{S}_e = \mathcal{S}.$$

Then \mathcal{S} is a finite subgraph of \mathcal{S} and we let \mathfrak{V} be the substructure of \mathfrak{S} induced on the set $V = V(\mathcal{S})$. The A -labelling of \mathcal{S} provides partial automorphisms of the structure \mathfrak{V} . Recall that the initialisation (3.17) of the σ_V -structure \mathfrak{S} on $V(\mathcal{S})$ actually took place in $V(\mathcal{S})$. This guarantees that an approximation of \mathfrak{S}^* is sitting in \mathfrak{V} . Let $\tilde{\mathfrak{V}}$ be an extension (on base set $\bar{V} \supseteq V$) of \mathfrak{V} on which all partial automorphisms of \mathfrak{V} given by \tilde{A} extend to total automorphisms. Let G be the corresponding A -generated group

of automorphisms of $\overline{\mathfrak{V}}$. The next result formulates and proves the main assertion about G .

Lemma 3.13. *If the coset labelling ℓ is G -commuting, then there exists a homomorphism $\mathfrak{C}^* \rightarrow \overline{\mathfrak{V}}$.*

Proof. For every edge $k \in K$ choose a word $w(k) \in \ell(k)$ such that the word labelling $w: K \rightarrow F$ is G -commuting. Let $k \in K$ and $w \in H_k g_k$; then in $\mathcal{S}_{H_k g_k}$ we have $\iota \cdot w = \tau$ where ι and τ are the initial and terminal vertices of $\mathcal{S}_{H_k g_k}$, respectively. Hence in $\overline{\mathfrak{V}}$ we have $\iota[w]_G = \tau$. This implies that $\varpi(k^{-1})[w(k)]_G = \varpi(k)$ also holds in $\overline{\mathfrak{V}}$. Now fix a vertex $o \in V$, and for every $v \in V$ choose a path $\pi_v: v \rightarrow o$ and let $w(\pi_v)$ be the label of that path. By assumption, the G -value $g_v := [w(\pi_v)]_G$ does not depend on the path π_v but only on the vertex v .

We modify the mapping $\varpi: K \rightarrow V$ to a mapping $\mu: K \rightarrow \overline{V}$ by setting, for $k \in K_v$, $\mu(k) := \varpi(k)g_v$ (note that $v = \omega k$). Clearly, μ is a homomorphism $\mathfrak{C} \rightarrow \overline{\mathfrak{V}}$ of σ_V -structures. Again let $k \in K$; the definition of $v \mapsto g_v$ implies $g_{\alpha k} = [w(k)]_G g_{\omega k}$ and we get

$$\begin{aligned} \mu(k^{-1}) &= \varpi(k^{-1})g_{\omega k^{-1}} = \varpi(k^{-1})g_{\alpha k} \\ &= \varpi(k^{-1})[w(k)]_G g_{\omega k} = \varpi(k)g_{\omega k} = \mu(k). \end{aligned}$$

It follows that $\equiv \subseteq \ker \mu$ and the homomorphism $\mu: \mathfrak{C} \rightarrow \overline{\mathfrak{V}}$ factors through \mathfrak{C}^* . \square

The proof of Proposition 3.11 is now immediate.

Proof of Proposition 3.11 by use of Theorem 3.2. Let the finite graph $\mathcal{E} = (V \cup K; \alpha, \omega, {}^{-1})$ be connected and without loop edges, and let ℓ be a coset labelling of \mathcal{E} which is not F -commuting. We need to show that ℓ is not G -commuting for some suitable finite A -generated group G .

According to the discussion after Proposition 3.11 it is sufficient to prove the statement for graphs without loop edges. From the graph \mathcal{E} we obtain the relational signature σ_V and an associated σ_V -structure \mathfrak{C} and its consolidated structure \mathfrak{C}^* . Let \mathbf{C} be the class $\mathbf{Excl}(\{\mathfrak{C}^*\})$ of all σ_V -structures \mathfrak{U} that do not admit a homomorphism $\mathfrak{C}^* \rightarrow \mathfrak{U}$. Now consider the structure \mathfrak{S} ; by Lemma 3.12 there is no homomorphism $\mathfrak{C}^* \rightarrow \mathfrak{S}$, that is, $\mathfrak{S} \in \mathbf{C}$. Let \mathfrak{V} be the (finite) substructure of \mathfrak{S} induced on the set $V = V(\mathfrak{S})$. Viewing the partial automorphisms of \mathfrak{V} that are induced by the A -graph structure of \mathfrak{S} as an extension problem $(\mathfrak{V}, \tilde{A})$, we have \mathfrak{S} as a solution in \mathbf{C} . By Theorem 3.2 there exists a finite solution $\overline{\mathfrak{V}}$ in \mathbf{C} . Let G be the A -generated group generated by the corresponding automorphisms. Since $\overline{\mathfrak{V}} \in \mathbf{C}$ there is no homomorphism $\mathfrak{C}^* \rightarrow \overline{\mathfrak{V}}$. By Lemma 3.13 then, the labelling ℓ is not G -commuting. \square

The argument that Theorem 3.3 follows from Proposition 3.11 is now similar to the second part of the proof of Theorem 3.5. There we looked at all n -tuples of elements of the inverse monoid M . An n -tuple may be

seen as a labelling of a cycle graph of length n . So let $\mathcal{E} = (V \cup K; \alpha, \omega, {}^{-1})$ be a finite connected graph. An M -labelling of \mathcal{E} is a mapping $l: K \rightarrow M$ such that $l(k^{-1}) = l(k)^{-1}$ for all $k \in K$. Every such labelling gives rise to a coset labelling ℓ by letting $\ell(k) = \varphi(l(k))$ where φ is defined in (2.3). By Proposition 2.10 every $\varphi(l(k))$ is a right coset $H_k g_k$ for some finitely generated subgroup H of F and such that $\varphi(l(k^{-1})) = (\varphi(l(k)))^{-1}$. Some of these coset labellings may be F -commuting. We take those labellings ℓ that are not F -commuting. For every such there exists a finite group G_ℓ such that ℓ is not G_ℓ -commuting. Then the A -generated direct product G over all such groups G_ℓ yields a finite group G required by Theorem 3.3.

Finally, we argue similarly as after the proof of Theorem 3.5 and show that, conversely, (Ash's) Theorem 3.3 implies Proposition 3.11. Let $\mathcal{E} = (V \cup K; \alpha, \omega, {}^{-1})$ be a finite graph and $\ell: K \rightarrow 2^F$, $k \mapsto H_k g_k$ a coset labelling not commuting with respect to F . For each geometric edge $e = \{k, k^{-1}\}$ of \mathcal{E} let \mathcal{S}_e be the plain Stallings graph of $H_k g_k$ and let $\mathcal{S} = \bigsqcup_{e \in K} \mathcal{S}_e$. Let $M = \mathcal{T}(\mathcal{S})$ and let G be a finite A -generated group as in Theorem 3.3 with respect to \mathcal{E} and M . We argue that ℓ is not G -commuting. Suppose, towards a contradiction, this were not the case and ℓ were G -commuting: there exists a G -commuting word labelling $l: K \rightarrow F$ covered by ℓ . According to Theorem 3.3 there exists an M -related and F -commuting relabelling l' of l . Let $k \in K$ and let $(\iota, \mathcal{S}_k, \tau)$ be the Stallings graph of $H_k g_k$; since $l(k) \in \ell(k) = H_k g_k$ we have $\iota \cdot [l(k)]_M = \tau$. Since l and l' are M -related we also have $\iota \cdot [l'(k)]_M = \tau$, which just means that $l'(k) \in H_k g_k = \ell(k)$. Since l' is F -commuting this contradicts the assumption that ℓ is not F -commuting.

3.5. Group-theoretic versions of the Herwig–Lascar Theorem. As mentioned in the introduction there is also a “group-theoretic version” of the Herwig–Lascar Theorem (Theorem 3.3 in [19]). There are even several variants of this version which we discuss briefly in this subsection; they deal with (non)solvability in the free group F of systems of equations of a certain type. Proposition 3.11 immediately gives us one of these versions, namely Proposition 3.14 below. However, in order to line up with the literature (Propositions 3.18 and 3.20 below) we have to modify the statement of Proposition 3.11 a bit: instead of labellings $\ell: k \mapsto H_k g_k$ of the edges of a graph \mathcal{E} by right cosets Hg we now consider labellings by left cosets gH . These labellings also have to satisfy $\ell(k^{-1}) = (\ell(k))^{-1}$ which now reads as $g_{k^{-1}} = g_k^{-1}$ and $H_{k^{-1}} = g_k^{-1} H_k^{g_k^{-1}}$ where $H^{g^{-1}} := gHg^{-1}$. So, from now on “coset labelling” means “left coset labelling”. It is clear that one version is as good as the other since every left coset is a right coset (with respect to some conjugate subgroup) and vice versa. The reason why left cosets occur in [19] where right cosets occur in our context is that Herwig–Lascar let (partial) automorphisms act on the left while we let them act on the right. We come to the first variant of the group-theoretic version of Herwig–Lascar, which is a consequence of the Proposition 3.11.

Proposition 3.14. *Let $\mathcal{E} = (V \cup K; \alpha, \omega, {}^{-1})$ be a finite graph, $\ell: K \rightarrow 2^F$, $k \mapsto g_k H_k$ a coset labelling and $(X_k)_{k \in K}$ variables. Then there exists a finite index normal subgroup N of F such that: if the system of equations*

$$X_{k_1} \cdots X_{k_t} = 1_F: k_1 \cdots k_t \text{ a closed path in } \mathcal{E} \quad (3.20)$$

has no solution in F subject to the constraints $X_k \in g_k H_k$, then it even has no solution subject to the constraints $X_k \in g_k H_k N$.

Lemma 3.15. *Proposition 3.11 and Proposition 3.14 are equivalent.*

Proof. Let \mathcal{E} be a finite graph with edge set K and let $k \mapsto g_k H_k$ be a (left) coset labelling. For all $k \in K$, let $x_k \in g_k H_k$ and $n_k \in N$. If for some $k_1, \dots, k_t \in K$, $x_{k_1} n_{k_1} \cdots x_{k_t} n_{k_t} = 1_F$ then $x_{k_1} \cdots x_{k_t} \in N$. This guarantees that Proposition 3.14 implies Proposition 3.11. For the converse assume that $(x_k)_{k \in K}$ is a G -commuting word labelling, i.e. that the choice of elements $x_k \in g_k H_k$ is such that for all closed paths $k_1 \cdots k_t$ in \mathcal{E} , the product $x_{k_1} \cdots x_{k_t}$ belongs to N . We need a choice of elements $n_k \in N$ for $k \in K$ such that for every such product $x_{k_1} \cdots x_{k_t}$,

$$x_{k_1} n_{k_1} \cdots x_{k_t} n_{k_t} = 1_F. \quad (3.21)$$

For every single such product a suitable choice is obvious. But we need a choice that works globally. For every connected component \mathcal{D} of \mathcal{E} choose a spanning tree $\mathcal{T}_{\mathcal{D}}$ (that is, a tree subgraph that connects all vertices of \mathcal{D}). For each edge $k \in \mathcal{T}_{\mathcal{D}}$ set $n_k = 1_F$. For $k \in \mathcal{D} \setminus \mathcal{T}_{\mathcal{D}}$ let $k_1 \cdots k_s$ be the unique reduced path $\omega k \rightarrow \alpha k$ in $\mathcal{T}_{\mathcal{D}}$. According to our assumption, $x_{k_1} \cdots x_{k_s} x_k =: m_k \in N$. Let $n_k := m_k^{-1}$; then, since $n_{k_i} = 1_F$ for the edges k_i ,

$$x_{k_1} n_{k_1} \cdots x_{k_s} n_{k_s} x_k n_k = x_{k_1} \cdots x_{k_s} x_k m_k^{-1} = 1_F$$

for all closed paths of the mentioned kind. By a standard argument (see e.g. the proof of Lemma 6.1 in [21]) it follows that for this choice of $(n_k)_{k \in K}$ the equality (3.21) holds for every closed path $k_1 \cdots k_t$ in \mathcal{E} . \square

By the standard argument from [21] mentioned in the preceding proof, the infinite system (3.20) is equivalent to some finite subsystem consisting of equations where the closed path $k_1 \cdots k_t$ runs through a set of reduced “basic cycles” whose cardinality is the cyclomatic number of \mathcal{E} , together with all “trivial cycles” kk^{-1} for $k \in K$.

Almeida and Delgado discussed the connection between Theorem 3.3 in [19] (i.e. the group-theoretic version of the Herwig–Lascar Theorem) and graph labellings in [4, 5] and showed that this theorem is equivalent to Ash’s Theorem for inverse monoids (Theorem 3.3 in the present paper). We need some further concepts for our discussion. Let $\mathcal{E} = (V \cup K; \alpha, \omega, {}^{-1})$ be a finite connected graph. An *extended word labelling* of \mathcal{E} is a mapping $w: V \cup K \rightarrow F$ such that $w \upharpoonright K$ is a word labelling. We also define *extended coset labelling* in the obvious way: it is a mapping $\ell: V \cup K \rightarrow 2^F$ such that $\ell(p) = g_p H_p$ for every $p \in V \cup K$ with all H_p finitely generated subgroups of F and $\ell \upharpoonright K$ is a coset labelling.

Let G be an A -generated group. An extended word labelling w of \mathcal{E} is G -coherent if for each edge $k \in K$, the equality

$$[w(\alpha k)]_G [w(k)]_G = [w(\omega k)]_G$$

holds. This simply means that $ab = c$ for all edges $k: \bullet \xrightarrow[a]{b} \bullet$ where $a = [w(\alpha k)]_G$, $b = [w(k)]_G$ and $c = [w(\omega k)]_G$. This notion is strongly connected with G -commutativity, as Lemmas 3.16 and 3.17 below show.

In the following, $\mathcal{E} = (V \cup K; \alpha, \omega, {}^{-1})$ is a finite connected graph.

Lemma 3.16. (1) Let $w: K \rightarrow F$ be a G -commuting word labelling of \mathcal{E} ; then for every $v \in V$ and $g \in G$ there exists a G -coherent extended word labelling \hat{w} of \mathcal{E} such that $[\hat{w}(v)]_G = g$ and $\hat{w} \upharpoonright K = w$.

(2) Conversely, for every G -coherent extended word labelling w of \mathcal{E} the restriction $w \upharpoonright K$ is G -commuting.

Proof. Let w be G -commuting; let $f \in F$ be a word such that $[f]_G = g$ and for every vertex $u \in V$ let $\pi_u: v \rightarrow u$ be a path from v to u . For $u \in V$ define $\hat{w}(u) := fw(\pi_u)$ and for $k \in K$ let $\hat{w}(k) := w(k)$. This gives an extended word labelling as required by (1). For (2) let w be a G -coherent extended word labelling. It follows by induction on the length that for every path $\pi: u \rightarrow v$, $[w(u)]_G [w(\pi)]_G = [w(v)]_G$, hence $[w(\pi)]_G = 1_G$ for every closed path π . \square

We use an extension of the given graph \mathcal{E} , which in [1, 4, 5] was called the *cone* of \mathcal{E}). Roughly speaking, we add to the graph \mathcal{E} a new vertex v_0 , and for every old vertex v a new edge $\vec{v}: v_0 \rightarrow v$ and its inverse \overleftarrow{v} . More formally, given $\mathcal{E} = (V \cup K; \alpha, \omega, {}^{-1})$ we let

$$\hat{V} := V \cup \{v_0\} \text{ where } v_0 \notin V$$

and

$$\hat{K} := K \cup \{\vec{v}, \overleftarrow{v} : v \in V\} \text{ where } \vec{v} \neq \overleftarrow{v} \text{ and } \vec{v}, \overleftarrow{v} \notin K$$

and set $\hat{\mathcal{E}} = (\hat{V} \cup \hat{K}; \alpha, \omega, {}^{-1})$ with incidence functions and inversion on the new edges \vec{v} and \overleftarrow{v} defined in the obvious way: $\alpha \vec{v} = v_0$, $\omega \vec{v} = v$ and $\vec{v}^{-1} = \overleftarrow{v}$ (this defines also $\alpha \overleftarrow{v}$ and $\omega \overleftarrow{v}$). The graph $\hat{\mathcal{E}}$ is designed to carry within a word labelling the information contained in an extended word labelling of \mathcal{E} . Indeed, given an extended word labelling $w: V \cup K \rightarrow F$ we define the word labelling \hat{w} of $\hat{\mathcal{E}}$ by $\hat{w}(k) = w(k)$ for every $k \in K$ and $\hat{w}(\vec{v}) = w(v)$ (and hence also $\hat{w}(\overleftarrow{v}) = w(v)^{-1}$). A proof of a statement equivalent to the following one forms a part of the proof of Theorem 7 in [1]; for the convenience of the reader we include a proof.

Lemma 3.17. An extended word labelling w of \mathcal{E} is G -coherent if and only if the word labelling \hat{w} of $\hat{\mathcal{E}}$ is G -commuting.

Proof. For the ‘if’ part it suffices to notice that for any edge k , the equality

$$[w(\alpha k)]_G [w(k)]_G = [w(\omega k)]_G$$

holds if and only if $[\widehat{w}(\overrightarrow{\alpha k} \cdot k \cdot \overleftarrow{\omega k})]_G = 1_G$. For the ‘only if’ part let $e_1 e_2 \cdots e_n$ be a closed path in $\widehat{\mathcal{E}}$. Immediately from the definition of the labelling \widehat{w} it follows that the claim is true for every trivial cycle path kk^{-1} with $k \in \widehat{K}$. By induction on the length of the path one can see that, in order to show that the G -value of its label is 1_G it is sufficient to assume that the path is simple, that is, with the exception of the initial and terminal vertices, no vertex is met twice. Also, instead of the path $e_1 \cdots e_n$ we may look at any shifted path $e_i \cdots e_n e_1 \cdots e_{i-1}$. Hence we may assume that the vertex v_0 is either the initial and the terminal vertex, or else does not occur at all. In the first case, the path is of the form $\overrightarrow{u_1} e_2 \cdots e_{n-1} \overleftarrow{u_{n-1}}$ for the vertices $u_1 = \alpha e_2$ and $u_{n-1} = \omega e_{n-1}$. Let u_2, \dots, u_{n-2} be the other vertices of the path $e_2 \cdots e_{n-1}$, that is $\omega e_i = u_i = \alpha e_{i+1}$. Then G -coherence of the labelling w implies that $[\widehat{w}(e_i)]_G = [\widehat{w}(\overleftarrow{u_{i-1}} \overrightarrow{u_i})]_G$ for every i . It follows that

$$[\widehat{w}(\overrightarrow{u_1} e_2 \cdots e_{n-1} \overleftarrow{u_{n-1}})]_G = [\widehat{w}(\overrightarrow{u_1} \overleftarrow{u_1} \overrightarrow{u_2} \cdots \overleftarrow{u_{n-2}} \overrightarrow{u_{n-1}} \overleftarrow{u_{n-1}})]_G = 1_G.$$

If v_0 does not occur on the path, let u_1, \dots, u_n be the vertices of the path, that is, $\omega e_n = u_1 = \alpha e_1$ and $\omega e_i = u_{i+1} = \alpha e_{i+1}$ for $i = 1, \dots, n-1$. Again by use of $[\widehat{w}(e_i)]_G = [\widehat{w}(\overleftarrow{u_i} \overrightarrow{u_{i+1}})]_G$ we get

$$[\widehat{w}(e_1 \cdots e_n)]_G = [\widehat{w}(\overleftarrow{u_1} \overrightarrow{u_2} \overleftarrow{u_2} \cdots \overleftarrow{u_n} \overrightarrow{u_1})]_G = [\widehat{w}(\overleftarrow{u_1} \overrightarrow{u_1})]_G = 1_G.$$

□

We now define an extended coset labelling $\ell: V \cup K \rightarrow 2^F$, $p \mapsto g_p H_p$ of a graph \mathcal{E} to be G -coherent for $G = F/N$ if there exists an extended word labelling $l: V \cup K \rightarrow F$ such that $l(p) \in \ell(p)$ for every $p \in V \cup K$ and l is G -coherent. This may be formulated as follows: for every $p \in V \cup K$ there exists a choice $x_p \in g_p H_p N$ such that the equality $x_{\alpha k} x_k = x_{\omega k}$ holds for all $k \in K$. This leads us to the Almeida–Delgado version of the (group-theoretic formulation of the) Herwig–Lascar Theorem [4, 5].

Proposition 3.18. *Let $\mathcal{E} = (V \cup K; \alpha, \omega, {}^{-1})$ be a finite graph and let $\ell: V \cup K \rightarrow 2^F$, $p \mapsto g_p H_p$ be an extended coset labelling and let $(X_p)_{p \in V \cup K}$ be variables. Then there exists a finite index normal subgroup N of F such that: if the system*

$$X_{\alpha k} X_k = X_{\omega k} : k \in K \tag{3.22}$$

has no solution in F subject to the constraints $X_p \in g_p H_p$ for all $p \in V \cup K$, then it even has no solution subject to the constraints $X_p \in g_p H_p N$ for all $p \in V \cup K$.

The equations in (3.22) do not require inversion and have the advantage that they make sense in situations other than groups. This led to the concepts of *hyperdecidability* and *tameness* of pseudovarieties of finite semi-groups and monoids [2, 3]. In fact, the formulation of the statement in [4, 5] does not require graphs with inversion of edges. The connections between G -commutativity and G -coherence discussed in Lemmas 3.16 and 3.17 show that Propositions 3.14 and 3.18 are equivalent. More precisely, the implication Proposition 3.14 \implies Proposition 3.18 follows from Lemma 3.17 while

the reverse implication Proposition 3.18 \implies Proposition 3.14 follows from Lemma 3.16.

Lemma 3.19. *Proposition 3.14 and Proposition 3.18 are equivalent.*

We now come to the formulation of the group-theoretic version of the Theorem of Herwig–Lascar [19, Theorem 3.3]. For finite index sets I and L consider tuples $(X_i)_{i \in I}$ of variables for elements from F , $(g_l)_{l \in L}$ of elements of F , and $(H_l)_{l \in L}$ of finitely generated subgroups of F . In their paper [19] Herwig and Lascar considered equations of the form $X_i \equiv X_j g_l \pmod{H_l}$ and $X_i \equiv g_l \pmod{H_l}$, respectively. Elements $x_i, x_j \in F$ solve these equations if and only if

$$x_i H_l = x_j g_l H_l \text{ and } x_i H_l = g_l H_l,$$

which is of course equivalent to $x_j^{-1} x_i \in g_l H_l$ and $x_i \in g_l H_l$, respectively. Hence these equations are equivalent to constraints of the form $X_j^{-1} X_i \in g_l H_l$ and $X_i \in g_l H_l$ the second of which occurs already in Propositions 3.14 and 3.18. Theorem 3.3 in [19] (i.e. the group-theoretic version of the Herwig–Lascar Theorem) then can be formulated as follows.

Proposition 3.20. *Let I and L be finite index sets, let $(X_i)_{i \in I}$, $(g_l)_{l \in L}$ and $(H_l)_{l \in L}$ be variables, elements of F and finitely generated subgroups of F and let $S \subseteq I \times L$ and $T \subseteq I^2 \times L$. Then there exists a finite index normal subgroup N of F such that: if the system of constraints*

$$X_i \in g_l H_l: (i, l) \in S \text{ and } X_j^{-1} X_i \in g_l H_l: (j, i, l) \in T \quad (3.23)$$

has no solution $(x_i) \in F^I$, then neither has the system

$$X_i \in g_l H_l N: (i, l) \in S \text{ and } X_j^{-1} X_i \in g_l H_l N: (j, i, l) \in T.$$

The Theorem of Ribes–Zalesskii can be expressed this way as follows: let H_1, \dots, H_n be finitely generated subgroups of F and $g \in F$. Consider variables X_1, \dots, X_{n-1} and constraints

$$X_{n-1} \in g H_n, X_{n-1}^{-1} X_{n-2} \in H_{n-1}, \dots, X_2^{-1} X_1 \in H_2, X_1 \in H_1.$$

One readily checks that this system has a solution if and only if $1_F \in g H_n \cdots H_1$, which is equivalent to $g \in H_1 \cdots H_n$. Hence the theorem implies that, if $g \notin H_1 \cdots H_n$, then there exists a finite index normal subgroup N of F such that $g \notin (H_1 N) \cdots (H_n N) = H_1 \cdots H_n N$, which implies $g N \cap H_1 \cdots H_n = \emptyset$. So, Proposition 3.20 can indeed be seen as a generalisation of the Ribes–Zalesskii Theorem.

We shall argue that Propositions 3.18 and 3.20 are equivalent. The easier direction is that Proposition 3.20 implies Proposition 3.18. Suppose we are given the situation described in the latter proposition. Solvability of the equation

$$X_{\alpha k} X_k = X_{\omega k}$$

subject to the constraints

$$X_{\alpha k} \in g_{\alpha k} H_{\alpha k}, X_k \in g_k H_k, X_{\omega k} \in g_{\omega k} H_{\omega k}$$

just means that the statement

$$\exists X_{\alpha k} \exists X_k \exists X_{\omega k} \left[\begin{array}{l} X_{\alpha k} \in g_{\alpha k} H_{\alpha k} \wedge X_k \in g_k H_k \\ \wedge X_{\omega k} \in g_{\omega k} H_{\omega k} \wedge X_{\alpha k} X_k = X_{\omega k} \end{array} \right] \quad (3.24)$$

is true. But (3.24) is equivalent to

$$\exists X_{\alpha k} \exists X_{\omega k} \left[X_{\alpha k} \in g_{\alpha k} H_{\alpha k} \wedge X_{\omega k} \in g_{\omega k} H_{\omega k} \wedge X_{\alpha k}^{-1} X_{\omega k} \in g_k H_k \right] \quad (3.25)$$

and the equivalence of these statements remains valid if every coset $g_l H_l$ involved is replaced by the coset $g_l H_l N$, for any fixed normal subgroup N of F . We see that the conditions occurring in (3.25) are just constraints as in Proposition 3.20. Hence, for every $k \in K$, (3.24) may be replaced by the equivalent (3.25). We therefore may convert the system of Proposition 3.18 into an equivalent one of the form in Proposition 3.20, and the equivalence remains valid if each coset $g_l H_l$ is replaced by the coset $g_l H_l N$, for any fixed (finite index) normal subgroup of F . Thus Proposition 3.20 implies Proposition 3.18.

For the converse we have to be more careful. Consider a finite system as in Proposition 3.20. We first look at the constraints $X_i \in g_l H_l$ in (3.23). If for some variable X_i no such constraint occurs, then we add the trivial one $X_i \in F$. Hence we may assume that every variable X_i occurs in at least one such constraint. So, for each $i \in I$ there exist group elements $g_{i1}, \dots, g_{in_i} \in F$ for some $n_i \geq 1$, and finitely generated subgroups H_{i1}, \dots, H_{in_i} of F such that the constraint $X_i \in g_l H_l$ occurs in (3.23) if and only if $g_l H_l = g_{is} H_{is}$ for some $s \in \{1, \dots, n_i\}$. For every $i \in I$ we introduce n_i new variables X_{i1}, \dots, X_{in_i} and consider a new system in the variables $V := \bigcup_{i \in I} \{X_{i1}, \dots, X_{in_i}\}$ defined as follows: instead of the constraints

$$X_i \in g_{i1} H_{i1}, \dots, X_i \in g_{in_i} H_{in_i}$$

take the constraints

$$\begin{aligned} &X_{i1} \in g_{i1} H_{i1}, \dots, X_{in_i} \in g_{in_i} H_{in_i} \text{ and } X_{is}^{-1} X_{it} \in \{1_F\} \\ &\text{for all } s, t = 1, \dots, n_i \text{ and } i \in I, \end{aligned} \quad (3.26)$$

and, for all $i, j \in I$, instead of every single constraint $X_j^{-1} X_i \in g_l H_l$ consider all constraints

$$X_{js}^{-1} X_{it} \in g_l H_l \text{ for all } s = 1, \dots, n_j \text{ and } t = 1, \dots, n_i. \quad (3.27)$$

The new system (3.26)&(3.27) is again a system of the same type as the one in (3.23). What is more, the old system (3.23) admits a solution if and only if the new system (3.26)&(3.27) admits one. Again this equivalence concerning solvability remains valid if each coset $g_l H_l$ involved is replaced by $g_l H_l N$, for any fixed normal subgroup N .

An important feature of the new system is that for every variable X there is exactly one constraint of the form $X \in gH$. We now proceed as in [4,

p. 415]⁴ and construct a “graph system” as in Proposition 3.18 that is equivalent to (3.26)&(3.27). We define a graph and an extended coset labelling ℓ as follows: the set of vertices is the set of variables of the system (3.26)&(3.27), and we set $\ell(X_{it}) = g_{it}H_{it}$ for every such vertex (i.e. variable); further, for two vertices (i.e. variables) $X_{it}, X_{js} \in V$ and every constraint $X_{js}^{-1}X_{it} \in g_lH_l$ we impose an edge $k: X_{js} \rightarrow X_{it}$ with label $\ell(k) = g_lH_l$; if necessary (that is, if the corresponding “inverse” constraint $X_{it}^{-1}X_{js} \in g_l^{-1}H_l^{g_l^{-1}}$ is not already contained in (3.27)) we add for every such edge its inverse edge $k^{-1}: X_{it} \rightarrow X_{js}$ and label it by $\ell(k^{-1}) = g_l^{-1}H_l^{g_l^{-1}}$. The graph and labelling so obtained provide a system of equations and a set of constraints as in Proposition 3.18. By the same arguments as discussed in connection with the equivalence of (3.24) and (3.25), we see that (3.26)&(3.27) and the graph system are equivalent, and this is still true if each constraint g_lH_l is replaced by g_lH_lN for some fixed normal subgroup N of F . Altogether, this means that, if we start with a system of the form (3.23) that has no solution, and if we assume Proposition 3.18 to be true, then we end up with the existence of a finite index normal subgroup N of F such that the system obtained from (3.23) by replacement of each coset g_lH_l by the coset g_lH_lN still has no solution. This just means that Proposition 3.18 implies Proposition 3.20. We have therefore proved the following.

Lemma 3.21. *Proposition 3.18 and Proposition 3.20 are equivalent.*

3.6. Construction of groups satisfying Ash’s Theorem. We briefly discuss the existence of and how to possibly construct a finite A -generated group G satisfying the conditions in Theorem 3.3 for a finite connected graph \mathcal{E} and a finite A -generated inverse monoid M . Let us call such a group G *suitable* for M and \mathcal{E} . Starting point is a finite A -generated group G that extends M in the sense of Definition 2.4. Such a group then is suitable for the given inverse monoid M and the cycle graph \mathcal{C}_1 of length 1 (that is, for the graph having one vertex and one geometric edge). Indeed, let $w: \{e, e^{-1}\} \rightarrow \tilde{A}^*$ be a G -commuting labelling; that is, $[w(e)]_G = 1_G$. Since G extends M , $[w(e)]_M$ is an idempotent and therefore $[w(e)]_M = [w(e)^{-1}]_M = [w(e)w(e)^{-1}]_M$; then $v(e) := w(e)w(e)^{-1}$ provides an M -related F -commuting labelling.

Throughout this subsection denote by \mathcal{C}_i the cycle graph of length i . We start with the construction of Ash. Based on $G_1 := G$ one can construct an expansion G_2 of G_1 that is suitable for M and \mathcal{C}_2 , an expansion G_3 of G_2 which is suitable for M and \mathcal{C}_3 , and so on, i.e. inductively an expansion G_k of G_{k-1} that is suitable for M and \mathcal{C}_k , for all k . Every group G_k is constructed in order to avoid relations of a certain kind which are satisfied by its predecessor G_{k-1} . The relations to be avoided by G_k can be described

⁴The argument in [4] does not take into account that some variables X_i in (3.23) may admit more than one constraint $X_i \in g_lH_l$.

in terms of the Cayley graph \mathcal{G}_{k-1} of the preceding group G_{k-1} . In order to make this more precise it is convenient to introduce the following concepts.

In the following we shall consider the indices in an n -tuple (x_1, \dots, x_n) as coming from a cyclic index set, that is, the indices in expressions like $x_{i\pm 1}$ have to be understood mod n ; this situation will be expressed by writing the tuple as $(x_i)_{i \in \mathbb{Z}_n}$. Let G be an A -generated group with Cayley graph \mathcal{G} . A 2-cyclic configuration $(\mathcal{P}_i, g_i)_{i \in \mathbb{Z}_2}$ of \mathcal{G} consists of a pair $\mathcal{P}_1, \mathcal{P}_2$ of connected subgraphs of \mathcal{G} and a pair g_1, g_2 of elements of G such that $g_2 = 1_G$ and g_1 and g_2 are contained in distinct connected components of $\mathcal{P}_1 \cap \mathcal{P}_2$. For $n \geq 3$, an n -cyclic configuration $(\mathcal{P}_i, g_i)_{i \in \mathbb{Z}_n}$ of \mathcal{G} consists of an n -tuple $\mathcal{P}_1, \dots, \mathcal{P}_n$ of connected subgraphs of \mathcal{G} and an n -tuple g_1, \dots, g_n of elements of G such that $g_n = 1_G$ and

- $g_i \in \mathcal{P}_i \cap \mathcal{P}_{i+1}$ for all i and
- $\mathcal{P}_{i-1} \cap \mathcal{P}_i \cap \mathcal{P}_{i+1} = \emptyset$ for all i .

The definition of n -cyclic configurations is reminiscent of the one of n -coset cycles (Definition 4.10) in Section 4; however, the context is considerably different. Figure 5 shows examples of 2-cyclic and 5-cyclic configurations.




FIGURE 4

An A -generated group H avoids the n -cyclic configuration $(\mathcal{P}_i, g_i)_{i \in \mathbb{Z}_n}$ of \mathcal{G} if the following holds:

- for every n -tuple of words u_1, \dots, u_n in \tilde{A}^* , if u_i labels a path from g_{i-1} to g_i in \mathcal{G} that runs entirely in \mathcal{P}_i , then $[u_1 \cdots u_n]_H \neq 1_H$.

In other words, generator sequences that follow the given cyclic configuration of the graphs \mathcal{P}_i in \mathcal{G} do not give rise to cycles in the Cayley graph \mathcal{H} of H .

A finite group H which avoids a given n -cyclic configuration can be constructed as follows. We form the disjoint union $\mathcal{P}_1 \sqcup \dots \sqcup \mathcal{P}_n$ of its graphs and factor by the smallest A -graph congruence that identifies the vertices $g_i^{(i)}$ of \mathcal{P}_i and $g_i^{(i+1)}$ of \mathcal{P}_{i+1} for $i = 1, \dots, n-1$, but not the vertices $g_n^{(n)}$ of \mathcal{P}_n and

$g_n^{(1)}$ of \mathcal{P}_1 . The result is a finite connected A -graph \mathcal{P} . The transition group of any finite completion $\bar{\mathcal{P}}$ of \mathcal{P} then avoids the given configuration. Since every finite \mathcal{G} has only finitely many n -cyclic configurations we can do this for each one; the A -generated direct product of the resulting groups then is a finite A -generated group that avoids all n -cyclic configurations of \mathcal{G} . The A -generated direct product of this group with the original group G then is an expansion H of G that avoids all n -cyclic configurations of \mathcal{G} .

We get a (non-deterministic) procedure that, given as input a finite A -generated inverse monoid M and a positive integer n , outputs a finite A -generated group G suitable for M and \mathcal{C}_n (recall that this means that G satisfies the condition of Theorem 3.3 with respect to M and \mathcal{C}_n). This procedure is as follows: starting with a group G_1 which is suitable for M and \mathcal{C}_1 (that is, G_1 extends M in the sense of Definition 2.4) we obtain a series of expansions $G_1 \leftarrow G_2 \leftarrow \cdots \leftarrow G_{n-1} \leftarrow G_n \leftarrow \cdots$ such that, in this series every G_k avoids all k -cyclic configurations of the Cayley graph \mathcal{G}_{k-1} of its predecessor G_{k-1} . Then, for every n , the n th group G_n in this series is suitable for M and \mathcal{C}_n (and hence for \mathcal{C}_m for all $m \leq n$). Let us call this procedure *n -cycle procedure*.

Next we obtain a procedure which on the input of a finite A -generated inverse monoid M and two positive integers k, n outputs a group G which is suitable for M and for all connected graphs \mathcal{E} of cyclomatic number k on n vertices. Indeed, we may iterate the n -cycle procedure k times. Suppose the inverse monoid M and the numbers n and k are given. We let G_1 be a group suitable for M and \mathcal{C}_n , $G_2 \twoheadrightarrow G_1$ a group suitable for $M \times G_1$ and \mathcal{C}_n , $G_3 \twoheadrightarrow G_2$ a group suitable for $M \times G_1 \times G_2$ and \mathcal{C}_n , and so on, $G_k \twoheadrightarrow G_{k-1}$ a group suitable for $M \times G_1 \times \cdots \times G_{k-1}$ and \mathcal{C}_n . Each group G_i is constructed by the application of the n -cycle procedure. By part 1 of the *proof of Lemma 3.7 by use of Theorem 3.5* it follows that every group G_i is critical for the pair $(M \times G_1 \times \cdots \times G_{i-1}, \mathcal{C}_n)$ in the sense of Definition 3.6. Recall that \times denotes the A -generated direct product. Then iterated application of (3.6) shows that G_k is critical for the pair (M, \mathcal{E}) and therefore is suitable for M and \mathcal{E} , provided that \mathcal{E} is connected on n vertices and has cyclomatic number k . In order to see this, we consider a chain $\mathcal{E}_0 \subseteq \mathcal{E}_1 \subseteq \cdots \subseteq \mathcal{E}_{k-1} \subseteq \mathcal{E}_k = \mathcal{E}$ where every \mathcal{E}_i has cyclomatic number i and \mathcal{E}_{i+1} is obtained from \mathcal{E}_i by adding an edge (and its inverse). (We note that if $H \twoheadrightarrow G$ then $G \times H \cong H$ for the A -generated direct product, so, in fact, $G_1 \times G_2 \times \cdots \times G_i \cong G_i$ for all i , but in the following it is useful to write the products in their redundant form.) Firstly we have that G_k is critical for $(M \times G_1 \times \cdots \times G_{k-1}, \mathcal{C}_n)$ and G_k is critical for $((M \times G_1 \times \cdots \times G_{k-1}) \times G_k, \mathcal{E}_0)$ (the latter holds for trivial reasons since \mathcal{E}_0 is a tree), hence by (3.6), G_k is critical for $(M \times G_1 \times \cdots \times G_{k-1}, \mathcal{E}_1)$. Next, G_{k-1} and therefore also G_k is critical for $(M \times G_1 \times \cdots \times G_{k-2}, \mathcal{C}_n)$ (since $G_k \twoheadrightarrow G_{k-1}$) and G_k is critical for $((M \times G_1 \times \cdots \times G_{k-2}) \times G_{k-1}, \mathcal{E}_1)$, so again by (3.6), G_k is critical for $(M \times G_1 \times \cdots \times G_{k-2}, \mathcal{E}_2)$. We can repeat this argument until we find that G_k is critical for (M, \mathcal{E}_k) and hence is suitable for M and $\mathcal{E}_k = \mathcal{E}$.

In order to perform the i th iteration of the above-mentioned procedure, that is, to find an expansion $G_i \twoheadrightarrow G_{i-1}$ that is critical for

$$(M \times G_1 \times \cdots \times G_{i-1}, \mathcal{C}_n),$$

it should be noticed that the group $G_{i-1} \cong G_1 \times \cdots \times G_{i-1}$ itself extends the inverse monoid $M \times G_1 \times \cdots \times G_{i-1}$, hence is the starting group for the application of the n -cycle procedure, which for the inverse monoid

$$M \times G_1 \times \cdots \times G_{i-1}$$

outputs a group critical for $(M \times G_1 \times \cdots \times G_{i-1}, \mathcal{C}_n)$. Altogether, this procedure for the construction of a group suitable for a given finite inverse monoid M and a given finite connected graph \mathcal{E} having n vertices and cyclomatic number k , boils down to the k -fold repetition of the n -cycle procedure.

From [8] it even follows that a group H suitable for M and any (fixed) finite connected graph \mathcal{E} is contained in any series of groups $G_1 \leftarrow G_2 \leftarrow \cdots$ starting with a group G_1 extending M and such that every group G_k in the series avoids all 2-cyclic configurations of the Cayley graph of its predecessor G_{k-1} . However, the proof is by a non-constructive topological compactness argument. In particular, it does not deliver a bound for a number t for which H can be chosen to be G_t . A concrete instance of such a series with an explicit bound is presented below.

The first author, partly in cooperation with B. Steinberg [7, 9, 10], also found a simpler construction for expansions of a group G which avoid cyclic configurations of \mathfrak{G} and therefore for the construction of groups which are suitable for a given finite inverse monoid M and a given finite connected graph \mathcal{E} . For an A -generated group G and an integer $p \geq 2$ there exists an A -generated expansion $\mathbf{Ab}_p G$ of G that is universal among all A -generated expansions K of G for which the kernel of $K \twoheadrightarrow G$ is an Abelian group of exponent dividing p : for every such expansion K , the canonical morphism $\mathbf{Ab}_p G \twoheadrightarrow G$ factors through K and the kernel of $\mathbf{Ab}_p G \twoheadrightarrow G$ itself is Abelian group of exponent dividing p . If $G \cong F/N$ then $\mathbf{Ab}_p G \cong F/[N, N]N^p$ and the group $\mathbf{Ab}_p G$ admits a transparent representation as a subgroup of a semidirect product of the free $|G \times A|$ -generated Abelian group of exponent p by G . It is shown in [9] that if G extends the monoid M , then $\mathbf{Ab}_p G$ is suitable for M and all graphs with two vertices (in semigroup-theoretic terms this says that $\mathbf{Ab}_p G$ is a witness for the group-pointlike sets of M). In particular, $\mathbf{Ab}_p G$ is suitable for M and the cycle graph \mathcal{C}_2 . We can iterate the expansion $G \mapsto \mathbf{Ab}_p G \mapsto \mathbf{Ab}_p(\mathbf{Ab}_p G) \mapsto \cdots$ and let $G^{(p,k)}$ denote the outcome of expanding G k times (with respect to p). It is shown in [7] that $G^{(p,n-1)}$ is suitable for M and \mathcal{C}_n for every n . We can further use the fact that loop edges (though they contribute to the cyclomatic number) do not contribute to the complexity concerned in the construction of suitable groups (see the proof of Theorem 3.3 by use of Lemma 3.7), and that only one expansion of the form $G \mapsto \mathbf{Ab}_p G$ is necessary to handle multiple edges. Combination of all this leads to the following

Proposition 3.22. *Let M be a finite A -generated inverse monoid and let G be a finite A -generated group extending M ; let \mathcal{E} be a connected graph on n vertices and let \mathcal{E}° be the graph obtained from \mathcal{E} by removing all loop edges and replacing every multiple edge by a single edge. Let k be the cyclomatic number of \mathcal{E}° . Then $G^{(p, (n-1)k+1)}$ is suitable for M and \mathcal{E}*

This is Theorem 4.3 in [7] except that for the number of expansions needed, instead of $(n-1)k+1$ the worse bound $(n-1)(2k-1)+1$ had been calculated there. In the sequence of iterations one may actually use different integers p_i at every iteration.

4. THEOREMS STRENGTHENED

The groups provided by the third author's construction [14, 27, 28] deserve special attention in semigroup theory. Indeed, they provide an affirmative solution to the Henckell–Rhodes Problem: every finite inverse monoid admits a finite F -inverse cover — see [14] for a detailed presentation of the topic. However, this is not the end of the story. By taking full advantage of the construction we get a strengthening of Ash's Theorem; in Section 5 some evidence is given that the new theorem is strictly stronger than Ash's Theorem. In turn, this leads to a strengthening of the Herwig–Lascar Theorem 3.2, see Theorem 4.18 below. We also get strengthenings of the group-theoretic versions of Section 3.5. The following is the intended strengthening of Ash's Theorem.

Theorem 4.1 (main theorem). *Let M be a finite A -generated inverse monoid and n a positive integer. Then there exists a finite A -generated group H such that, for every finite connected graph \mathcal{E} on at most n vertices, every H -commuting labelling v of \mathcal{E} admits a labelling u such that*

- (1) u commutes over the free group F ,
- (2) $[v(e)]_M = [u(e)]_M$ for every edge $e \in \mathcal{E}$,
- (3) $[v(e)]_H = [u(e)]_H$ for every edge $e \in \mathcal{E}$.

In short, every H -commuting labelling of \mathcal{E} admits an $M \times H$ -related relabelling that commutes over the free group F . This theorem relates to Ash's Theorem just as the construction of a finite A -generated F -inverse cover of a finite A -generated inverse monoid M relates to the construction of a finite group witnessing the group-pointlike sets of M (see [14, Section 2.6]). Note that for a group H satisfying Theorem 4.1 the inverse monoid $M \times H$ is an F -inverse cover of M (provided that \mathcal{E} contains a cycle graph of length at least 2). Theorem 4.1 is proved in Section 4.2. In preparation for that poof we recall some key concepts and results from [14, 27, 28] in Section 4.1.

4.1. Summary of some results from [14, 27, 28]. We use the notation of [14]. In this section, we shall be concerned with two different generating sets: A and E . The set E denotes the set of positive edges of some finite

connected oriented graph \mathcal{Q} (to be specified later). In Section 4.1.1 we treat the general case where E is the set of positive edges of a finite connected oriented graph \mathcal{Q} while in Section 4.1.2, $E = Q \times A$ is the set of positive edges of the Cayley graph \mathcal{Q} of some finite A -generated group Q .

4.1.1. *The general case.* For $a \in E$ and $p \in \tilde{E}^*$ let $p_{a \rightarrow 1}$ be the word obtained from p by deletion of all occurrences of a and a^{-1} in p . An E -generated group G is *retractable* if, for all words $p, q \in \tilde{E}^*$ and every letter $a \in E$ the following holds:⁵

$$[p]_G = [q]_G \implies [p_{a \rightarrow 1}]_G = [q_{a \rightarrow 1}]_G.$$

Definition 4.2. For a word $p \in \tilde{E}^*$, the *word content* $\text{co}(p)$ is the set of all letters $a \in E$ for which a or a^{-1} occurs in p .

Assume that G is retractable: for $p, q \in \tilde{E}^*$ and $a \in E$ the equality $[p]_G = [q]_G$ implies $[p_{a \rightarrow 1}]_G = [q_{a \rightarrow 1}]_G$. Suppose now that $a \in \text{co}(p)$ but $a \notin \text{co}(q)$. Then the words q and $q_{a \rightarrow 1}$ are identical. Hence $[p]_G = [q]_G$ implies

$$[p_{a \rightarrow 1}]_G = [q_{a \rightarrow 1}]_G = [q]_G = [p]_G.$$

In this way, we may (without changing the value $[p]_G$) delete every letter in a word p that does not occur in every other word representation q of the group element $[p]_G$. This leads to the following definition.

Definition 4.3. Let G be a retractable E -generated group and $g \in G$. The *content* $C(g)$ of g is

$$C(g) := \bigcap \{ \text{co}(q) : q \in \tilde{E}^*, [q]_G = g \}.$$

For a word $p \in \tilde{E}^*$ the G -*content* of p is the content $C([p]_G)$.

Remark 4.4. For any retractable E -group G and any word $w \in \tilde{E}^*$ the word content $\text{co}(w)$ of w always contains the G -content $C([w]_G)$ of w .

For a retractable E -generated group G , the G -content of a word $p \in \tilde{E}^*$ is the unique \subseteq -minimal subset B of E such that $[p]_G$ can be represented as a word over \tilde{B} . For a given word p denote by p° the word obtained from p by deleting every occurrence of a letter (and its inverse) of p not in the G -content of p . Then p° is a representation with minimal content of the group element $[p]_G$, that is, $[p^\circ]_G = [p]_G$ and $\text{co}(p^\circ) = C([p^\circ]_G) = C([p]_G)$.

Now let G be some retractable E -generated group and let $p \in \tilde{E}^*$ be a word forming a path $p: u \rightarrow v$ in the graph \mathcal{Q} (recall that E is the set of positive edges of the oriented graph \mathcal{Q}). If we form the aforementioned minimal representation p° of $[p]_G$ as a word over \tilde{E} , deletion of letters in p in general destroys the *path quality* of p with respect to \mathcal{Q} : p° is *some* minimal representation of $[p]_G$, but we cannot expect that p° forms a path in \mathcal{Q} , let alone a path $u \rightarrow v$.

⁵It suffices to restrict this postulate to the case $q = 1$.

The main result of [14, Sections 3–5] describes the construction of a *finite* E -generated group G having the following crucial property [14, Cor. 5.7].

Lemma 4.5. *For every word $p \in \tilde{E}^*$ forming a path $u \rightarrow v$ in \mathcal{Q} , there exists a word $q \in \tilde{E}^*$ such that*

- (1) q also forms a path $u \rightarrow v$ in \mathcal{Q} ,
- (2) $[q]_G = [p]_G$,
- (3) $\text{co}(q) = \text{co}(p^\circ)$.

That is, q is G -related to p , uses only letters (i.e. edges) from the common G -content and also forms a path $u \rightarrow v$ in \mathcal{Q} . Note that finiteness of G is of the essence here, since the free E -generated group obviously satisfies these conditions.

The group G is constructed in a series of expansions

$$G_1 \leftarrow H_1 \leftarrow G_2 \leftarrow \cdots \leftarrow G_{|E|-1} \leftarrow H_{|E|-1} \leftarrow G_{|E|} = G, \quad (4.1)$$

the members of which are the transition groups of complete E -graphs

$$\mathcal{X}_1 \subseteq \mathcal{Y}_1 \subseteq \mathcal{X}_2 \subseteq \cdots \subseteq \mathcal{X}_{|E|-1} \subseteq \mathcal{Y}_{|E|-1} \subseteq \mathcal{X}_{|E|}, \quad (4.2)$$

that is, $G_i = \mathcal{T}(\mathcal{X}_i)$ and $H_i = \mathcal{T}(\mathcal{Y}_i)$. Every graph in the series (4.2) is obtained from its predecessor by adding certain complete components tailored from the Cayley graph of the corresponding preceding group; the (complicated) details are worked out in [14, Sections 3–5]. The graph \mathcal{X}_1 is a certain completion of the original graph \mathcal{Q} that is considered as an E -graph in which every edge is labelled by itself (observe that the labelling alphabet E coincides with the set of positive edges E of \mathcal{Q}). In addition, the completion \mathcal{X}_1 of \mathcal{Q} is obtained by adding just edges so that \mathcal{Q} and \mathcal{X}_1 have the same set of vertices. Let us choose and fix an arbitrary vertex of \mathcal{Q} (i.e. of \mathcal{X}_1) and denote it by $1_{\mathcal{Q}}$. From the construction it follows that for every group G_i and H_i in the series (4.1) there exists a unique canonical graph morphism $\varphi_i: \mathcal{G}_i \rightarrow \mathcal{X}_1$ (respectively $\psi_i: \mathcal{H}_i \rightarrow \mathcal{X}_1$) mapping 1_{G_i} (respectively 1_{H_i}) to $1_{\mathcal{Q}}$ where \mathcal{G}_i (respectively \mathcal{H}_i) is the Cayley graph of G_i (respectively H_i). To simplify notation put $H := H_{|E|-1}$ and recall that $G := G_{|E|}$, and denote their Cayley graphs as \mathcal{G} respectively \mathcal{H} . We let

$$\widehat{\mathcal{Q}} := \text{the connected component of } 1_G \text{ in } \varphi_{|E|}^{-1}(\mathcal{Q}) \quad (4.3)$$

and likewise

$$\widehat{\mathcal{Q}}_H := \text{the connected component of } 1_H \text{ in } \psi_{|E|-1}^{-1}(\mathcal{Q})$$

and call these graphs the \mathcal{Q} -cover at 1_G in \mathcal{G} and the \mathcal{Q} -cover at 1_H in \mathcal{H} , respectively. For the canonical mapping $\phi: \mathcal{G} \rightarrow \mathcal{H}$ induced by the canonical morphism $G \rightarrow H$ we have $\phi(\widehat{\mathcal{Q}}) = \widehat{\mathcal{Q}}_H$.

Restricting the canonical graph morphism $\varphi_{|E|}: \mathcal{G} \rightarrow \mathcal{X}_1$ to $\widehat{\mathcal{Q}}$, we get the canonical graph morphism $\widehat{\mathcal{Q}} \rightarrow \mathcal{Q}$ whose restriction to the vertex set we denote as $\gamma \mapsto \bar{\gamma}$. For the next lemma, a path Π in an E -graph is a B -path for some subset $B \subseteq E$ if the label of Π is a word over \tilde{B} . The crucial

property of the content function C of G , mentioned in Lemma 4.5, implies the following result.

Lemma 4.6. *Let γ, δ be vertices of $\widehat{\mathcal{Q}}$ and $\Pi: \gamma \rightarrow \delta$ be a B -path in \mathcal{G} for some $B \subseteq E$; then there exists a B -path $\Lambda: \gamma \rightarrow \delta$ which runs in $\widehat{\mathcal{Q}}$.*

Proof. Recall that the label $\ell(_)$ of every path in the following is a word over \widetilde{E} , E the generating set of G . Let $P: \gamma \rightarrow \delta$ be some path in $\widehat{\mathcal{Q}}$ that links the start and end points of the given $\Pi: \gamma \rightarrow \delta$. Then $\gamma[\ell(\Pi)]_G = \delta = \gamma[\ell(P)]_G$, whence $[\ell(\Pi)]_G = [\ell(P)]_G$. In particular, $\ell(\Pi)$ and $\ell(P)$ have the same G -content, which is necessarily contained in B (since the word content of $\ell(\Pi)$ is contained in B). Now $\ell(P)$ forms a path $\bar{\gamma} \rightarrow \bar{\delta}$ in \mathcal{Q} . By Lemma 4.5, there exists a word ϱ over \widetilde{E} which forms a path $\bar{\gamma} \rightarrow \bar{\delta}$ in \mathcal{Q} such that $[\ell(P)]_G = [\varrho]_G$ and such that ϱ uses only letters (i.e. edges) from the G -content of $\ell(P)$, which is contained in B . The path $\varrho: \bar{\gamma} \rightarrow \bar{\delta}$ admits a unique lift $\widehat{\varrho}: \gamma \rightarrow \delta'$ in $\widehat{\mathcal{Q}}$ for some $\delta' \in \mathcal{Q}$ for which $\bar{\delta} = \bar{\delta}'$. Note that the label $\ell(\widehat{\varrho})$ of $\widehat{\varrho}$ is just ϱ . Since $[\ell(P)]_G = [\varrho]_G$ we also have $\delta = \gamma[\ell(P)]_G = \gamma[\varrho]_G = \gamma[\ell(\widehat{\varrho})]_G$; hence $\widehat{\varrho}: \gamma \rightarrow \delta$ and this path runs in $\widehat{\mathcal{Q}}$. Since this path uses only letter (i.e. edges) from its G -content, it is a B -path as required. \square

We fix some notation. For an E -graph \mathcal{E} , a subset $B \subseteq E$ and a vertex q of \mathcal{E} , $q\mathcal{E}[B]$ denotes the B -component of \mathcal{E} containing q : the subgraph of \mathcal{E} spanned by all B -paths in \mathcal{E} starting at q . If $\varphi: \mathcal{E} \rightarrow \mathcal{D}$ is a canonical morphism of complete E -graphs, q a vertex of \mathcal{E} and $B \subsetneq E$ then $\varphi(q\mathcal{E}[B]) = \varphi(q)\mathcal{D}[B]$. For the Cayley graph \mathcal{G} of G (being an E -graph), some $\gamma \in G$ and subset $K \subseteq E$, the graph $\gamma\mathcal{G}[K]$ is the K -component of \mathcal{G} at γ ; we shall refer to such graphs as K -coset graphs or just *coset graphs* if the set K is clear from the context. As a K -graph, this is clearly isomorphic to the Cayley graph of the subgroup $G[K]$ of G generated by K . Of crucial use will be Lemma 4.9 below (being an extension of Lemma 4.6): it follows from the construction of G in [14] but it is neither mentioned nor used in that paper (but it is a key ingredient in [28]: see Definition 6.13 and Lemma 6.17 there). To this end we present some more details of the construction of G in [14, Section 5]. This concerns a specific component \mathcal{U} of $\mathcal{X}_{|E|}$ (as earlier, we let $H = H_{|E|-1}$ and \mathcal{H} be the Cayley graph of H).

Lemma 4.7. *The graph $\mathcal{X}_{|E|}$ contains a connected component \mathcal{U} satisfying the following conditions:*

- (1) \mathcal{U} contains $\widehat{\mathcal{Q}}_H$ as a subgraph,
- (2) for all proper subsets $B, C \subsetneq E$ and all vertices $\xi, \eta \in \widehat{\mathcal{Q}}_H \subseteq \mathcal{U}$ we have:
 - (a) $\xi\mathcal{U}[B] \cap \eta\mathcal{U}[C] \neq \emptyset \implies \xi\mathcal{U}[B] \cap \eta\mathcal{U}[C] \cap \widehat{\mathcal{Q}}_H \neq \emptyset$,
 - (b) $\xi\mathcal{U}[B] \cap \eta\mathcal{U}[C] = \kappa\mathcal{U}[B \cap C]$ for any $\kappa \in \xi\mathcal{U}[B] \cap \eta\mathcal{U}[C]$,
 - (c) the graphs $\xi\mathcal{U}[B] \cap \widehat{\mathcal{Q}}_H$ and $\eta\mathcal{U}[C] \cap \widehat{\mathcal{Q}}_H$ are connected.

Lemma 4.8. *Let \mathcal{U} be the graph of Lemma 4.7 and $\vartheta \in \widehat{\mathcal{Q}} \subseteq \mathcal{G}$; then for any $D \subsetneq E$, the canonical morphism $\varphi: \mathcal{G} \rightarrow \mathcal{U}$ (mapping 1_G to $1_H \in \widehat{\mathcal{Q}}_H$) is injective in restriction to the coset subgraph $\vartheta\mathcal{G}[D]$.*

For proofs of these assertions, the reader is referred to Sections 3.3.2, 4 and 5.2, in particular, Definition 3.16, (3.8), Theorem 4.6 and the proof of Condition 5.2, (ii) and (iii) (for $k = |E|$) in [14]. In the notation of [14], the graph \mathcal{U} occurs as $\overline{\text{CE}}(H, \widehat{\mathcal{Q}}_H; \mathbb{P}_E)$ in item (2) of Definition 5.3 for $k = |E| - 1$, $H_k = H$, $A = E$, $\mathcal{E} = \mathcal{A} = \langle E \rangle = \mathcal{Q}$ and $\mathcal{E}_{\mathcal{H}_k} = \mathcal{Q}_{\mathcal{H}} = \widehat{\mathcal{Q}}_H$.

Lemma 4.9. *Let $K, L \subsetneq E$ and let γ, δ be vertices of $\widehat{\mathcal{Q}}$. If $\gamma\mathcal{G}[K] \cap \delta\mathcal{G}[L] \neq \emptyset$ in \mathcal{G} , then $\gamma\mathcal{G}[K] \cap \delta\mathcal{G}[L] \cap \widehat{\mathcal{Q}} \neq \emptyset$. Moreover, if ε is a vertex in that intersection, then there exist a K -path $\gamma \rightarrow \varepsilon$ and an L -path $\delta \rightarrow \varepsilon$, both of which run in $\widehat{\mathcal{Q}}$.*

Proof. Let $\theta \in \gamma\mathcal{G}[K] \cap \delta\mathcal{G}[L]$ and consider the canonical graph morphism $\varphi: \mathcal{G} \rightarrow \mathcal{U}$ (mapping 1_G to $1_H \in \widehat{\mathcal{Q}}_H$). According to Lemma 4.8 the restrictions $\varphi_\gamma := \varphi \upharpoonright \gamma\mathcal{G}[K]$ and $\varphi_\delta := \varphi \upharpoonright \delta\mathcal{G}[L]$ are injective. Hence $\varphi_\gamma: \gamma\mathcal{G}[K] \rightarrow \varphi(\gamma)\mathcal{U}[K]$ and $\varphi_\delta: \delta\mathcal{G}[L] \rightarrow \varphi(\delta)\mathcal{U}[L]$ are graph isomorphisms. Moreover, $\varphi(\theta) \in \varphi(\gamma)\mathcal{U}[K] \cap \varphi(\delta)\mathcal{U}[L]$. In particular, $\varphi(\gamma)\mathcal{U}[K] \cap \varphi(\delta)\mathcal{U}[L] \neq \emptyset$. Since $\varphi(\gamma), \varphi(\delta) \in \varphi(\widehat{\mathcal{Q}}) = \widehat{\mathcal{Q}}_H$ it follows from Lemma 4.7 (2a) that $\varphi(\gamma)\mathcal{U}[K] \cap \varphi(\delta)\mathcal{U}[L] \cap \widehat{\mathcal{Q}}_H \neq \emptyset$. Let ξ be a vertex in this intersection; by Lemma 4.7 (2b) there exists a $(K \cap L)$ -path $\Upsilon: \varphi(\theta) \rightarrow \xi$ (running in $\varphi(\theta)\mathcal{U}[K \cap L]$). Since φ_γ and φ_δ agree on $\gamma\mathcal{G}[K] \cap \delta\mathcal{G}[L] = \theta\mathcal{G}[K \cap L]$ (this equality holds by retractability of G), their inverses φ_γ^{-1} and φ_δ^{-1} agree on $\varphi(\theta)\mathcal{U}[K \cap L]$. We then have $\varphi_\gamma^{-1}(\Upsilon) = \varphi_\delta^{-1}(\Upsilon)$ which is a $(K \cap L)$ -path $\theta \rightarrow \hat{\xi} := \varphi_\gamma^{-1}(\xi) = \varphi_\delta^{-1}(\xi)$ in $\theta\mathcal{G}[K \cap L] = \gamma\mathcal{G}[K] \cap \delta\mathcal{G}[L]$. Since $\varphi(\gamma)\mathcal{U}[K] \cap \widehat{\mathcal{Q}}_H$ is connected by Lemma 4.7 (2c), so is $\varphi_\gamma^{-1}(\varphi(\gamma)\mathcal{U}[K] \cap \widehat{\mathcal{Q}}_H)$. It follows that $\gamma = \varphi_\gamma^{-1}(\varphi(\gamma))$ and $\hat{\xi} = \varphi_\gamma^{-1}(\xi)$ are in the same connected component of $\varphi^{-1}(\widehat{\mathcal{Q}}_H)$, that is $\hat{\xi} \in \widehat{\mathcal{Q}}$. Consequently, $\gamma\mathcal{G}[K] \cap \delta\mathcal{G}[L] \cap \widehat{\mathcal{Q}} \neq \emptyset$. The second assertion of the lemma follows immediately from Lemma 4.6. \square

In addition to the previous facts we require a further property of G , namely *n-acyclicity* for some positive integer n (to be specified later) which we are going to discuss now. The notions of n -coset cycles and n -acyclic groups are at the heart of [27, 28]. In the following we use again the cyclic index set \mathbb{Z}_n , that is, indices $i \in \mathbb{Z}_n$ are understood mod n .

Definition 4.10. Let G be an E -generated retractable group, $B \subseteq E$, $E_1, \dots, E_n \subsetneq B$ and $\eta_1, \dots, \eta_n \in G$. The subgraph

$$\eta_1\mathcal{G}[E_1] \cup \eta_2\mathcal{G}[E_2] \cup \dots \cup \eta_n\mathcal{G}[E_n]$$

of \mathcal{G} is an *n-coset cycle* provided that:

- (1) $\eta_{i+1} \in \eta_i\mathcal{G}[E_i]$ for all $i \in \mathbb{Z}_n$,
- (2) $\eta_i\mathcal{G}[E_{i-1} \cap E_i] \cap \eta_{i+1}\mathcal{G}[E_i \cap E_{i+1}] = \emptyset$ for all $i \in \mathbb{Z}_n$.

This notion is reminiscent of the notion of an n -cyclic configuration in Section 3.6. In fact, n -coset cycles are special types of n -cyclic configurations. The context, however, is somehow different. In Section 3.6 we had a group G and we wanted to construct an expansion H of G that avoids a given n -cyclic configuration (or even all of them) *of its quotient* G . In the present context we are looking for a group H which avoids all n -coset cycles *in its own* Cayley graph, which is a more ambitious task.

Definition 4.11. Let $n > 1$ be a positive integer; an E -generated retractable group G is n -acyclic if the Cayley graph \mathcal{G} does not contain any coset cycles of length up to n .

Note that 2-acyclicity of G is the same as $G[B \cap C] = G[B] \cap G[C]$ for all subsets $B, C \subsetneq E$. This just says that two coset subgraphs $g\mathcal{G}[B]$ and $h\mathcal{G}[C]$ of \mathcal{G} either have empty intersection or otherwise $g\mathcal{G}[B] \cap h\mathcal{G}[C] = k\mathcal{G}[B \cap C]$ for any $k \in g\mathcal{G}[B] \cap h\mathcal{G}[C]$; this condition is guaranteed by retractability of G . We briefly sketch how we can achieve that the group G in question be n -acyclic for any given (but fixed) positive integer n . Let $B \subseteq E$, $B_1, \dots, B_n \subsetneq B$ and let H be an E -generated group with Cayley graph \mathcal{H} ; for every $i \in \mathbb{Z}_n$ choose a vertex $\gamma_i \in \mathcal{H}[B_i]$ such that

$$\mathcal{H}[B_{i-1} \cap B_i] \cap \gamma_i \mathcal{H}[B_i \cap B_{i+1}] = \emptyset.$$

Note that $\mathcal{H}[B_{i-1} \cap B_i] \cup \gamma_i \mathcal{H}[B_i \cap B_{i+1}]$ is a subgraph of $\mathcal{H}[B_i]$. Now take the disjoint union $\bigsqcup_{i=1}^n \mathcal{H}[B_i]$ and factor by the smallest E -graph congruence Θ which identifies $\gamma_i \in \mathcal{H}[B_i]$ with $1 \in \mathcal{H}[B_{i+1}]$, $1 \leq i \leq n-1$. Call the resulting graph a ϖB amalgamation chain of length n for the group H . One can check that for $C \subsetneq B$, every C -component of a ϖB amalgamation chain of length n is either a C -coset $\gamma \mathcal{H}[C]$ or a ϖC amalgamation chain of length at most n (see [27, 28]). It follows that for $n \geq 6$, n -acyclicity of G is achieved by adding in Definition 5.3 of the graph \mathcal{Z}_k in [14] as a third item (3) all ϖA -amalgamation chains of length up to $n-2$ for the group H_k in the series (4.1) for all subsets $A \subseteq E$ of size $|A| = k+1$ (for details see [27, 28]).⁶

4.1.2. *The case of \mathcal{Q} being the Cayley graph of an A -generated group Q .* We now assume that the oriented graph \mathcal{Q} is the Cayley graph of a finite A -generated group Q for some finite alphabet A . Recall that the set of vertices of \mathcal{Q} is Q and the set of positive edges E of \mathcal{Q} is $Q \times A$ where

$$\alpha(q, a) = q, \quad \omega(q, a) = q[a]_Q \text{ and } (q, a)^{-1} = (q[a]_Q, a^{-1}). \quad (4.4)$$

The last equality in (4.4) implies the equality of sets $(Q \times A)^{-1} = Q \times A^{-1}$ and hence $\tilde{E} = Q \times \tilde{A}$ (this is strictly true and not just a manner of speaking).

Let G be the E -generated group as discussed in the previous subsection, but for \mathcal{Q} the Cayley graph of a finite A -generated group Q . We shall use

⁶The group G as it is defined in [14] is already 5-acyclic since the augmented clusters in item (1) of Definition 5.3 contain all amalgamation chains of length up to 3.

the Cayley graph \mathcal{G} of G , which has vertex set G and edge set $G \times \tilde{E}$. The edges of \mathcal{G} therefore are of the form $(\gamma, (q, a))$ with $\gamma \in G$, $(q, a) \in \tilde{E}$ and

$$(\gamma, (q, a))^{-1} = (\gamma[(q, a)]_G, (q[a]_Q, a^{-1})).$$

Recall (4.3), the definition of $\hat{\mathcal{Q}}$, the \mathcal{Q} -cover at 1_G in \mathcal{G} (in the previous subsection 1_Q was a vertex of \mathcal{Q} chosen arbitrarily, now we let 1_Q be the identity element of the group Q). The vertex set of $\hat{\mathcal{Q}}$ is denoted \hat{Q} . The surjective mapping $\hat{Q} \twoheadrightarrow Q$ induced by the canonical morphism $\hat{\mathcal{Q}} \rightarrow \mathcal{Q}$ will be denoted $\gamma \mapsto \bar{\gamma}$ (as in the previous subsection).

Remark 4.12. The (canonical) mapping $\hat{Q} \twoheadrightarrow Q$ can be extended to a surjective mapping $G \rightarrow Q$ by restricting the canonical graph morphism $\mathcal{G} \rightarrow \mathcal{X}_1$ to its vertex set and taking into account that the vertex set of \mathcal{X}_1 is Q . This mapping $G \rightarrow Q$ neither is a group homomorphism nor does it give rise to a graph homomorphism $\mathcal{G} \rightarrow \mathcal{Q}$, and in this paper we will use only its restriction to \hat{Q} .

In the following we shall deal with A -graphs (related to subgraphs of \mathcal{Q}) as well as E -graphs (related to subgraphs of \mathcal{G}). The E -graph $\hat{\mathcal{Q}}$ is the subgraph of \mathcal{G} spanned by all paths in \mathcal{G} starting at 1_G and whose labels (being words over \tilde{E}) form paths in \mathcal{Q} starting at 1_Q . Every such path in \mathcal{Q} is *induced* by a word over \tilde{A} : for every word $x \in \tilde{A}^*$ and every $q \in Q$ there exists a unique path in \mathcal{Q} with initial vertex q whose label is x — we denote this path by $\pi_q^{\mathcal{Q}}(x)$ and say this path is *induced* by the word x . We give a short summary of the situation.

Summary 4.13. Let

$$\Pi = (\gamma_1, (q_1, a_1)) \cdots (\gamma_n, (q_n, a_n))$$

be a word over $G \times \tilde{E}$. Then Π forms a path $\gamma \rightarrow \delta$ in $\hat{\mathcal{Q}}$ if and only if its label

$$\ell(\Pi) = (q_1, a_1) \cdots (q_n, a_n)$$

(a word over \tilde{E}) forms a path $\bar{\gamma} \rightarrow \bar{\delta}$ in \mathcal{Q} induced by the word $x = a_1 \dots a_n$, that is

$$\ell(\Pi) = \pi_{\bar{\gamma}}^{\mathcal{Q}}(x).$$

Moreover, in this case, $\gamma = \gamma_1$ and $\bar{\gamma} = \bar{\gamma}_1 = q_1$. The canonical graph morphism $\hat{\mathcal{Q}} \rightarrow \mathcal{Q}$ maps the path Π onto the path $\ell(\Pi) = \pi_{\bar{\gamma}}^{\mathcal{Q}}(x)$. The original path Π can be seen as the path in \mathcal{G} starting at γ induced by the word $\ell(\Pi)$, and this path actually runs in $\hat{\mathcal{Q}} \subseteq \mathcal{G}$.

By [14, Proposition 5.5] the construction of G guarantees that every bijection φ of E that is induced by an automorphism of the oriented graph \mathcal{Q} extends to an automorphism $\hat{\varphi}$ of G . In particular, this is the case for every automorphism φ_q of \mathcal{Q} induced by multiplication by the element q on the right. Hence the permutation $\varphi_q: E \rightarrow E$ (for $E = Q \times A$), $(g, a) \mapsto {}^q(g, a) := (qg, a)$ extends to an automorphism of G . It follows

that the group Q acts on the group G by automorphisms on the left via $\gamma \mapsto {}^q\gamma$ for $\gamma \in G, q \in Q$. So we can form the semidirect product $G \rtimes Q$.

In [14] the group H (depending on G and Q) is defined as:

$$H := \langle [(1_Q, a)]_G, [a]_Q : a \in A \rangle \leq G \rtimes Q, \quad (4.5)$$

which is an A -generated group. (Compare Section 2.5 in [14] and especially equation (2.3) there.) For a word x over \tilde{A} , we have $[x]_H = ([\pi_1^Q(x)]_G, [x]_Q)$ (see (2.4) in [14]).

In a sense, the group H is sitting inside G . More precisely, the vertices of $\hat{\mathcal{Q}}$ are in bijective correspondence with the elements of H via the map $\gamma \mapsto (\gamma, \bar{\gamma})$. Indeed, take any word $x \in \tilde{A}^*$; then the terminal vertex of the path $\pi_1^Q(x)$ is $[x]_Q$. Being a word over \tilde{E} , $\pi_1^Q(x)$ induces a path Π in \mathcal{G} starting at 1_G running in $\hat{\mathcal{Q}}$. Its terminal vertex is $[\ell(\Pi)]_G = [\pi_1^Q(x)]_G$. The canonical graph morphism $\hat{\mathcal{Q}} \twoheadrightarrow \mathcal{Q}$ maps Π onto $\ell(\Pi) = \pi_1^Q(x)$, and hence maps the terminal vertex $[\ell(\Pi)]_G = [\pi_1^Q(x)]_G$ of Π onto the terminal vertex $[x]_Q$ of $\pi_1^Q(x)$. Consequently, the element $([\pi_1^Q(x)]_G, [x]_Q)$ is of the form $(\gamma, \bar{\gamma})$ for $\gamma = [\pi_1^Q(x)]_G$, that is, the mapping $\hat{\mathcal{Q}} \rightarrow H, \gamma \mapsto (\gamma, \bar{\gamma})$ is surjective. Clearly, for two distinct vertices $\gamma, \eta \in \hat{\mathcal{Q}}$ we have $(\gamma, \bar{\gamma}) \neq (\eta, \bar{\eta})$. As a consequence we have the following.

Lemma 4.14. *For two words $x, y \in \tilde{A}^*$, $[x]_H = [y]_H$ if and only if $[\pi_q^Q(x)]_G = [\pi_q^Q(y)]_G$ for some (any) $q \in Q$.*

This result also shows that the group H could equally well be defined as the quotient $H := F/N$ where F is the A -generated free group and

$$N = \{[w]_F : w \in \tilde{A}^*, [\pi_1^Q(w)]_G = 1_G\}.$$

Without going too much into the details, we mention that the “equivalence” of $\hat{\mathcal{Q}}$ and H admits a groupoidal interpretation. Since the generating set of the group G is the set E of positive edges of the graph \mathcal{Q} we can as well consider the *groupoid* \mathcal{Q} generated by the graph \mathcal{Q} subject to the relations of G defined as follows: the vertices (i.e. objects) of \mathcal{Q} are the vertices of \mathcal{Q} , that is the elements of Q . For $p, q \in Q$ the arrows (i.e. morphisms) $p \rightarrow q$ are all G -values of paths $p \rightarrow q$ in \mathcal{Q} (recall that every such path is a word over \tilde{E}):

$$\mathcal{Q}(p, q) = \{[\pi]_G : \pi \text{ is a path } p \rightarrow q \text{ in } \mathcal{Q}\}.$$

The G -value $[\pi]_G$ of every path π in \mathcal{Q} starting at 1_Q is the terminal vertex $\omega\hat{\pi}$ of the lift $\hat{\pi}$ of π in \mathcal{G} starting at 1_G . Hence the set

$$\mathcal{H} := \bigsqcup_{q \in Q} \mathcal{Q}(1_Q, q) \quad (4.6)$$

of arrows of \mathcal{Q} with initial vertex 1_Q is in bijective correspondence with the set of vertices of $\hat{\mathcal{Q}}$ that is, with the elements of $\hat{\mathcal{Q}}$. We can define a group

structure on this set of arrows of \mathcal{Q} (and hence on \widehat{Q}) as follows. For two such arrows $\gamma \in \mathcal{Q}(1_Q, \bar{\gamma})$, $\delta \in \mathcal{Q}(1_Q, \bar{\delta})$ we define

$$\gamma\delta := \gamma \cdot \bar{\gamma}\delta$$

where the multiplication \cdot on the right hand side is the multiplication in the groupoid \mathcal{Q} and $\bar{\gamma}\delta$ is the arrow $\bar{\gamma} \longrightarrow \bar{\gamma}\bar{\delta}$ in \mathcal{Q} obtained by “shifting” the arrow δ from $1_Q \longrightarrow \bar{\delta}$ to $\bar{\gamma} \longrightarrow \bar{\gamma}\bar{\delta}$. To make this more precise, recall that Q acts on its Cayley graph \mathcal{Q} by left multiplication: $g \mapsto {}^r g := rg$ and $(g, a) \mapsto {}^r(g, a) := (rg, a)$ for $r \in Q$, g a vertex of \mathcal{Q} and (g, a) an edge of \mathcal{Q} . This action extends to an action by groupoid automorphisms of \mathcal{Q} in a natural way: on the set of vertices (i.e. objects) the action is the same as for \mathcal{Q} ; on the set of arrows $[\pi]_G: p \longrightarrow q$ it becomes ${}^r[\pi]_G: rp \longrightarrow rq$ where ${}^r[\pi]_G = [{}^r\pi]_G$ and ${}^r\pi = {}^r e_1 \cdots {}^r e_n$ for $\pi = e_1 \cdots e_n$ with $e_i \in \tilde{E}$. So the arrow $\bar{\gamma}\delta: \bar{\gamma} \longrightarrow \bar{\gamma}\bar{\delta}$ is just the result of the application of the automorphism $x \mapsto \bar{\gamma}x$ to the arrow $\delta: 1_Q \longrightarrow \bar{\delta}$. The resulting group is A -generated, namely by (the G -values of) the edges $1_Q \xrightarrow{a} [a]_Q$, more precisely by the set $\{[(1_Q, a)]_G: a \in A\}$, and is isomorphic to the group H .

In the notation of Margolis and Pin [24], the group so obtained is exactly \mathcal{Q}/Q , see Proposition 3.11 in [24]. Here \mathcal{Q}/Q denotes the set of orbits of the action of Q on \mathcal{Q} on which the groupoid structure of \mathcal{Q} induces a (quotient) groupoid structure which turns out to be a group since the transitivity of the action of Q on the set of vertices (i.e. objects) of \mathcal{Q} forces \mathcal{Q}/Q to have only one vertex (i.e. one object). The set \mathcal{H} in 4.6 is a complete set of representatives of the orbits, hence may be used to define the quotient structure. Altogether, the groupoidal interpretation just discussed may be summarised in one sentence: *ignore the (redundant) second components $\bar{\gamma}$ in the elements $(\gamma, \bar{\gamma}) \in H$, but compose the first components γ correctly.*

4.2. Proof of Theorem 4.1. The task of this section is to prove Theorem 4.1. Lemma 4.15 is the main achievement in this section. While inverse monoids are seemingly absent in the lemma, the connection with inverse monoids will become clear in the statements after its proof.

We let Q be a finite A -generated group with Cayley graph \mathcal{Q} and consider the A -generated group H defined in (4.5), where G is the E -generated group, for $E = Q \times A$, as discussed in Sections 4.1.1 and 4.1.2. In particular G satisfies the conditions of Lemmas 4.5 and 4.9 and we also assume that G is n -acyclic in the sense of Definition 4.11 for some positive integer n — the existence of such a group has been outlined in the last two paragraphs of Section 4.1.1. Note that G and H are uniquely determined by the input data Q and n according to the construction outlined in Sections 4.1.1 and 4.1.2 and we set $H = H(Q, n)$. Recall the word content $\text{co}(_)$ of a word over \tilde{E} according to Definition 4.2 and that, for a word w over \tilde{A} and $g \in Q$, $\pi_1^Q(w)$ and $\pi_g^Q(w)$ denote the paths in \mathcal{Q} labelled w , starting at 1_Q and at g respectively, considered as words over \tilde{E} .

Lemma 4.15 (main lemma). *Let Q be a finite A -generated group, n a positive integer and $H = H(Q, n)$. Then for any finite connected graph $\mathcal{E} = (V \cup K; \alpha, \omega, {}^{-1})$ on at most n vertices and without loop edges and for any H -commuting word labelling $v: K \rightarrow \tilde{A}^*$ there exist a finite tree $\mathcal{T} = (W \cup L; \alpha, \omega, {}^{-1})$ with vertex set $W \supseteq V$ and a word labelling $w: L \rightarrow \tilde{A}^*$ such that: for every edge $k \in K$, if π_k is the reduced path $\alpha k \rightarrow \omega k$ in \mathcal{T} and $w_k = w(\pi_k)$, then*

- (1) $[v(k)]_H = [w_k]_H$ and
- (2) $\text{co}(\pi_1^Q(v(k))) \supseteq \text{co}(\pi_1^Q(w_k))$.

In addition, \mathcal{T} can be chosen such that all vertices of $W \setminus V$ have degree at least 3.

Proof. The proof of this lemma is split into several cases, reflecting increasing levels of complexity with respect to the graph structure, in § 4.2.1– 4.2.4. In the following, elements of Q will always be denoted by g, h, g_i, h_i , etc., those of G by $\gamma, \eta, \gamma_i, \eta_i$, etc., and those of H by $(\gamma, g), (\gamma_i, g_i)$, etc.

4.2.1. Graphs with two vertices. Let us consider first all graphs having only two vertices, say o_1, o_2 and $t > 1$ geometric edges between these vertices, and let $k_1, \dots, k_t: o_1 \rightarrow o_2$. This case includes cycle graphs of length 2. Let $v_i := v(k_i)$; H -commutativity of the labelling just says that $[v_1]_H = \dots = [v_t]_H$, that is,

$$([\pi_1^Q(v_i)]_G, [v_i]_Q) = [v_i]_H = [v_j]_H = ([\pi_1^Q(v_j)]_G, [v_j]_Q)$$

for all i, j . In particular, the Q -values of all words v_i coincide, say $q := [v_i]_Q$ for all i . Also, the G -values of all sequences $\pi_1^Q(v_i)$ coincide, say $\gamma := [\pi_1^Q(v_i)]_G$ for all i . A fortiori, all sequences $\pi_1^Q(v_i)$ have the same G -content, say $E_0 := C(\gamma)$; also note that $E_0 \subseteq \text{co}(\pi_1^Q(v_i))$ for all i (the word content always contains the G -content). According to Lemma 4.5 there exists a word π over \tilde{E}_0 forming a path $\pi: 1_Q \rightarrow q$ in \mathcal{Q} such that $[\pi]_G = \gamma$. Since π forms a path in \mathcal{Q} it is induced by a word w over \tilde{A} , that is, $\pi = \pi_1^Q(w)$ for some word $w \in \tilde{A}^*$. Then $\text{co}(\pi_1^Q(w)) = E_0 \subseteq \text{co}(\pi_1^Q(v_i))$ for all i and $[\pi_1^Q(w)]_G = \gamma$. From Lemma 4.14 we get $[w]_H = [v_1]_H = \dots = [v_t]_H$. A suitable tree \mathcal{T} for the given graph labelling consists of the two vertices o_1, o_2 , an edge $f: o_1 \rightarrow o_2$ and its inverse f^{-1} , and the labelling $w(f) := w$.

4.2.2. Cycle graphs of length 3. We continue with cycle graphs of length 3. Let $\mathcal{C} = (V \cup K; \alpha, \omega, {}^{-1})$ be a cycle graph of length 3; let $V = \{o_1, o_2, o_3\}$ and $K = \{k_1^{\pm 1}, k_2^{\pm 1}, k_3^{\pm 1}\}$ with $k_i: o_i \rightarrow o_{i+1}$ (indices to be taken mod 3). Let $v: K \rightarrow \tilde{A}^*$ be an H -commuting labelling, given by $v(k_i) = v_i$. For $i = 1, 2, 3$, let $[v_i]_H =: (\gamma_i, g_i)$ so that $g_i = [v_i]_Q$ and $\gamma_i = [\pi_1^Q(v_i)]_G$. We set h_i and η_i according to

$$\left. \begin{array}{ll} h_1 := 1_Q & \eta_1 := 1_G \\ h_2 := g_1 & \eta_2 := \gamma_1 \\ h_3 := g_1 g_2 & \eta_3 := \gamma_1 \gamma_2 \end{array} \right\} \text{ so that } \left\{ \begin{array}{l} (\eta_i, h_i) = [v_1 \cdots v_{i-1}]_H \\ \text{and } h_i = \bar{\eta}_i \text{ for } 1 \leq i \leq 3. \end{array} \right.$$




FIGURE 5

For $i = 1, 2, 3$ let $E_i \subseteq E$ be the content of ${}^{h_i}\gamma_i$, which in fact is the G -content of $\pi_{h_i}^Q(v_i)$. Since G is 3-acyclic, the subgraph

$$\eta_1\mathcal{G}[E_1] \cup \eta_2\mathcal{G}[E_2] \cup \eta_3\mathcal{G}[E_3]$$

of \mathcal{G} is not a coset cycle. Since $\eta_{i+1} = \eta_i {}^{h_i}\gamma_i \in \eta_i\mathcal{G}[E_i]$ (indices to be taken mod 3, that is, $1_G = \eta_1 = \eta_3 {}^{h_3}\gamma_3 \in \eta_3\mathcal{G}[E_3]$), 3-acyclicity implies that

$$\eta_1\mathcal{G}[E_1] \cap \eta_2\mathcal{G}[E_2] \cap \eta_3\mathcal{G}[E_3] \neq \emptyset. \quad (4.7)$$

Since $\eta_1 \in \eta_3\mathcal{G}[E_3]$ we have $\eta_3\mathcal{G}[E_3] = \eta_1\mathcal{G}[E_3]$, and (4.7) can be written as

$$\eta_1\mathcal{G}[E_1 \cap E_3] \cap \eta_2\mathcal{G}[E_2] = \eta_1\mathcal{G}[E_1] \cap \eta_2\mathcal{G}[E_2] \cap \eta_1\mathcal{G}[E_3] \neq \emptyset. \quad (4.8)$$

Here we use that retractability of G implies $G[E_1] \cap G[E_3] = G[E_1 \cap E_3]$. Let \hat{Q} be the cover of Q at 1_G in \mathcal{G} . Since $1_G = \eta_1$ and η_2 belong to \hat{Q} we have by Lemma 4.9 that

$$\eta_1\mathcal{G}[E_1 \cap E_3] \cap \eta_2\mathcal{G}[E_2] \cap \hat{Q} \neq \emptyset.$$

Let δ be some vertex in this intersection, see Figure 5.

By Lemma 4.9 there are paths Π_i running in \hat{Q} , $\Pi_i: \eta_i \rightarrow \delta$, such that Π_1 runs in $\eta_1\mathcal{G}[E_1 \cap E_3]$, Π_2 runs in $\eta_2\mathcal{G}[E_1 \cap E_2]$ and Π_3 runs in $\eta_3\mathcal{G}[E_2 \cap E_3]$. Since the paths Π_i run in \hat{Q} they are induced by words over \tilde{A} : there are $x_1, x_2, x_3 \in \tilde{A}^*$ such that $\ell(\Pi_i) = \pi_{h_i}^Q(x_i)$. From the definition of the paths

Π_i it follows that

$$\begin{aligned} h_1 \gamma_1 &= [\ell(\Pi_1 \Pi_2^{-1})]_G = [\pi_{h_1}^Q(x_1 x_2^{-1})]_G \\ h_2 \gamma_2 &= [\ell(\Pi_2 \Pi_3^{-1})]_G = [\pi_{h_2}^Q(x_2 x_3^{-1})]_G \\ h_3 \gamma_3 &= [\ell(\Pi_3 \Pi_1^{-1})]_G = [\pi_{h_3}^Q(x_3 x_1^{-1})]_G \end{aligned}$$

where $h_1 = 1_Q$. Figure 5 depicts the situation: the shaded area represents the Q -cover \tilde{Q} and the dotted arrows indicate the paths with labels $\pi_{h_i}^Q(v_i)$. From Lemma 4.14 it follows that

$$[v_1]_H = [x_1 x_2^{-1}]_H, [v_2]_H = [x_2 x_3^{-1}]_H, [v_3]_H = [x_3 x_1^{-1}]_H.$$

Finally, from the definition of E_i ,

$$\begin{aligned} \text{co}(\pi_{h_1}^Q(x_1 x_2^{-1})) &= E_1 \subseteq \text{co}(\pi_{h_1}^Q(v_1)) \\ \text{co}(\pi_{h_2}^Q(x_2 x_3^{-1})) &= E_2 \subseteq \text{co}(\pi_{h_2}^Q(v_2)) \\ \text{co}(\pi_{h_3}^Q(x_3 x_1^{-1})) &= E_3 \subseteq \text{co}(\pi_{h_3}^Q(v_3)) \end{aligned}$$

where $\text{co}(_)$ denotes the word content of the sequence in question (being a word over $\tilde{E} = Q \times \tilde{A}$). Also note that, for any words y, z over \tilde{A} and any $g \in Q$, $\pi_1^Q(y) \subseteq \pi_1^Q(z)$ if and only if $\pi_g^Q(y) \subseteq \pi_g^Q(z)$.

We construct the required tree $\mathcal{T} = (W \cup L; \alpha, \omega, ^{-1})$ as follows: we let $W = V \cup \{p\} = \{o_1, o_2, o_3, p\}$ with $p \neq o_i$ for $i = 1, 2, 3$ and edges $L = \{f_1^{\pm 1}, f_2^{\pm 1}, f_3^{\pm 1}\}$ where $f_i: o_i \rightarrow p$ and labelling $w: L \rightarrow \tilde{A}^*$ given by $w(f_i) = x_i$. Then, for all edges $k \in \mathcal{E}$, if $\pi_k: \alpha k \rightarrow \omega k$ is the unique reduced path in \mathcal{T} and $w_k = w(\pi_k)$, then $[v(k)]_H = [w_k]_H$ and $\text{co}(\pi_1^Q(w_k)) \subseteq \text{co}(\pi_1^Q(v(k)))$, as required. Figure 5 indicates how the cycle (dotted arrows) is converted into a tree (solid arrows with additional vertex δ), though this happens in the G -component of the corresponding H -values.

From Figure 5 one might get the impression that we have tacitly assumed that $\delta \neq \eta_i$ and $[v_i]_H \neq 1_H$ for all i . However it may happen that $\delta = \eta_i$ for some i ; in this case the path Π_i and the word x_i reduce to the empty path and empty word, respectively. The label $w(f_i)$ of the arrow f_i will in turn be the empty word 1. Since we allow the labelling of some edges of \mathcal{T} to be the empty word 1, this is not in conflict with our construction. Similarly, if for some i , $[v_i]_H = 1_H$, then two of the vertices η_j in Figure 5 would coincide. This would have the effect that two of the words x_j coincide, which again would not be in conflict with our construction.

4.2.3. Cycle graphs of arbitrary length. We proceed by induction to prove the claim for all cycle graphs. Let $n > 3$ and assume the claim for all cycle graphs of length at most $n - 1$. In the following we use the cyclic index set \mathbb{Z}_n . Let $\mathcal{C} = (V \cup K; \alpha, \omega, ^{-1})$ be a cycle graph of length n with vertex set $V = \{o_1, \dots, o_n\}$ and let $k_i: o_i \rightarrow o_{i+1}$, $i \in \mathbb{Z}_n$, be its edges.

Let an H -commuting labelling $v: K \rightarrow \tilde{A}^*$ be given by $v(k_i^{\pm 1}) = v_i^{\pm 1}$ and let $[v_i]_H = (\gamma_i, g_i)$, that is, $g_i = [v_i]_Q$ and $\gamma_i = [\pi_1^Q(v_i)]_G$. For $i = 0, \dots, n - 1$

we set $h_{i+1} = g_1 \cdots g_i$ and $\eta_{i+1} = {}^{h_1}\gamma_1 \cdots {}^{h_i}\gamma_i$ (note that $h_1 = 1_Q$ and $\eta_1 = 1_G$). Then $[v_1 \cdots v_{i-1}]_H = (\eta_i, h_i)$ and in particular $\bar{\eta}_i = h_i$ (for $i = 1, \dots, n$). Let E_i be the content of ${}^{h_i}\gamma_i = [\pi_{h_i}^Q(v_i)]_G$, that is, the G -content of $\pi_{h_i}^Q(v_i)$ (see Definitions 4.2 and 4.3). Since G is n -acyclic, the subgraph

$$\eta_1 \mathcal{G}[E_1] \cup \eta_2 \mathcal{G}[E_2] \cup \cdots \cup \eta_n \mathcal{G}[E_n]$$

of \mathcal{G} is not a coset cycle. Since $\eta_{i+1} = \eta_i {}^{h_i}\gamma_i \in \eta_i \mathcal{G}[E_i]$ and $1_G = \eta_1 = \eta_n {}^{h_n}\gamma_n \in \eta_n \mathcal{G}[E_n]$, condition (2) in Definition 4.10 of coset cycles must be violated: for some i

$$\eta_i \mathcal{G}[E_{i-1} \cap E_i] \cap \eta_{i+1} \mathcal{G}[E_i \cap E_{i+1}] \neq \emptyset.$$




FIGURE 6

All elements $1_G = \eta_1, \eta_2, \dots, \eta_n$ belong to \hat{Q} . From Lemma 4.9 we get

$$\eta_i \mathcal{G}[E_{i-1} \cap E_i] \cap \eta_{i+1} \mathcal{G}[E_i \cap E_{i+1}] \cap \hat{Q} \neq \emptyset.$$

Let δ be a vertex in this intersection; then, in particular,

$$\delta \in \eta_{i-1} \mathcal{G}[E_{i-1}] \cap \eta_{i+1} \mathcal{G}[E_{i+1}] \cap \hat{Q}.$$

Hence there is an E_{i-1} -path $\Pi_1: \eta_{i-1} \longrightarrow \delta$ and an E_{i+1} -path $\Pi_2: \eta_{i+2} \longrightarrow \delta$, both running in \hat{Q} . Likewise there is an $(E_{i-1} \cap E_i)$ -path $\Lambda_1: \eta_i \longrightarrow \delta$ and an $(E_i \cap E_{i+1})$ -path $\Lambda_2: \eta_{i+1} \longrightarrow \delta$, both running in \hat{Q} . The situation is depicted in Figure 6, where the dotted arrows indicate the paths with labels $\pi_{h_j}^Q(v_j)$ for $j = i-1, i, i+1$, the shaded area the Q -cover \hat{Q} . Note that $\eta_j \mathcal{G}[E_j] = \eta_{j+1} \mathcal{G}[E_j]$ since $\eta_{j+1} \in \eta_j \mathcal{G}[E_j]$ for $j = i-1, i, i+1$.

Since these four paths run in $\widehat{\mathcal{Q}}$ they are induced by words in \widetilde{A}^* . So, there are $x_1, x_2, y_1, y_2 \in \widetilde{A}^*$ such that

$$\ell(\Pi_1) = \pi_{h_{i-1}}^{\mathcal{Q}}(x_1), \ell(\Pi_2) = \pi_{h_{i+2}}^{\mathcal{Q}}(x_2), \ell(\Lambda_1) = \pi_{h_i}^{\mathcal{Q}}(y_1), \ell(\Lambda_2) = \pi_{h_{i+1}}^{\mathcal{Q}}(y_2).$$

This leads to the following equalities

$$\begin{aligned} [\pi_{h_{i-1}}^{\mathcal{Q}}(v_{i-1})]_G &= [\ell(\Pi_1 \Lambda_1^{-1})]_G = [\pi_{h_{i-1}}^{\mathcal{Q}}(x_1 y_1^{-1})]_G \\ [\pi_{h_i}^{\mathcal{Q}}(v_i)]_G &= [\ell(\Lambda_1 \Lambda_2^{-1})]_G = [\pi_{h_i}^{\mathcal{Q}}(y_1 y_2^{-1})]_G \\ [\pi_{h_{i+1}}^{\mathcal{Q}}(v_{i+1})]_G &= [\ell(\Lambda_2 \Pi_2^{-1})]_G = [\pi_{h_{i+1}}^{\mathcal{Q}}(y_2 x_2^{-1})]_G \end{aligned}$$

so that, by Lemma 4.14,

$$[v_{i-1}]_H = [x_1 y_1^{-1}]_H, [v_i]_H = [y_1 y_2^{-1}]_H \text{ and } [v_{i+1}]_H = [y_2 x_2^{-1}]_H, \quad (4.9)$$

which implies

$$[v_{i-1} v_i v_{i+1}]_H = [x_1 x_2^{-1}]_H$$

and therefore also

$$[v_1 \cdots v_{i-2} x_1 x_2^{-1} v_{i+1} \cdots v_n]_H = 1_H. \quad (4.10)$$

Furthermore, by construction, the paths $\Pi_1 \Lambda_1^{-1}$, $\Lambda_1 \Lambda_2^{-1}$ and $\Lambda_2 \Pi_2^{-1}$ run in $\eta_{i-1} \mathcal{G}[E_{i-1}]$, $\eta_i \mathcal{G}[E_i]$ and $\eta_{i+1} \mathcal{G}[E_{i+1}]$, respectively. For the respective contents of their labels this implies

$$\text{co}(\pi_{h_{i-1}}^{\mathcal{Q}}(v_{i-1})) \supseteq E_{i-1} = \text{co}(\pi_{h_{i-1}}^{\mathcal{Q}}(x_1 y_1^{-1})), \quad (4.11)$$

$$\text{co}(\pi_{h_i}^{\mathcal{Q}}(v_i)) \supseteq E_i = \text{co}(\pi_{h_i}^{\mathcal{Q}}(y_1 y_2^{-1})), \quad (4.12)$$

$$\text{co}(\pi_{h_{i+1}}^{\mathcal{Q}}(v_{i+1})) \supseteq E_{i+1} = \text{co}(\pi_{h_{i+1}}^{\mathcal{Q}}(y_2 x_2^{-1})) \quad (4.13)$$

since the G -content of a word is always contained in its word content.

In order to reduce the issue to the inductive hypothesis, we consider a cycle graph \mathcal{C}_0 of length $n - 1$ with sets of vertices

$$V_0 = \{o_1, \dots, o_{i-1}, p, o_{i+2}, \dots, o_n\}$$

and of edges

$$\{k_t : t = 1, \dots, i-2, i+2, \dots, n\} \cup \{f_1, f_2\}$$

(and their respective inverses), where $f_1 : o_{i-1} \rightarrow p$ and $f_2 : p \rightarrow o_{i+2}$ and the edges k_t are as in \mathcal{C} , namely $k_t : o_t \rightarrow o_{t+1}$ for all $t \neq i-1, i, i+1$. We define a labelling v_0 of \mathcal{C}_0 by setting $v_0(k_t) = v_t$ (v_t as for \mathcal{C} , for all $t \neq i-1, i, i+1$) and $v_0(f_1) = x_1$ and $v_0(f_2) = x_2^{-1}$. In other words, in order to get \mathcal{C}_0 from \mathcal{C} , the labelled arc spanned by

$$o_{i-1} \xrightarrow{v_{i-1}} o_i \xrightarrow{v_i} o_{i+1} \xrightarrow{v_{i+1}} o_{i+2}$$

in \mathcal{C} is replaced by the labelled arc spanned by

$$o_{i-1} \xrightarrow{x_1} p \xrightarrow{x_2^{-1}} o_{i+2}.$$

According to (4.10), the labelling v_0 commutes over H . Hence by the inductive assumption there exists a finite tree $\mathcal{T}_0 = (W_0 \cup L_0, \alpha, \omega, {}^{-1})$ with

$V_0 \subseteq W_0$ such that all vertices in $W_0 \setminus V_0$ have degree at least 3. We may assume that $o_i, o_{i+1} \notin W_0$. By the inductive hypothesis, there is a labelling $w_0: L_0 \rightarrow \tilde{A}^*$ such that, for every edge k of \mathcal{C}_0 , if $\pi_k: \alpha k \rightarrow \omega k$ is the unique reduced path in \mathcal{T}_0 , then $[v_0(k)]_H = [w_0(\pi_k)]_H$ and $\text{co}(\pi_1^Q(w_0(\pi_k))) \subseteq \text{co}(\pi_1^Q(v_0(k)))$.




FIGURE 7. The tree \mathcal{T}

We define the tree $\mathcal{T} = (W \cup L, \alpha, \omega, ^{-1})$ according to the picture in Figure 7: let $W = W_0 \cup \{o_i, o_{i+1}\}$ and $L = L_0 \cup \{l_1^{\pm 1}, l_2^{\pm 1}\}$ where $l_1: o_i \rightarrow p$ and $l_2: o_{i+1} \rightarrow p$. We have $V \subseteq W$; the vertex p is not in V and has degree at least 3 in \mathcal{T} by construction. All other vertices not in V are in $W_0 \setminus V_0$ and have degree at least 3 by the inductive assumption. We define the labelling $w: L \rightarrow \tilde{A}^*$ by $w(l_t^{\pm 1}) = y_t^{\pm 1}$ ($t = 1, 2$) and $w(l) = w_0(l)$ for every edge $l \in L_0$. Let $\pi_{f_1}: o_{i-1} \rightarrow p$ and $\pi_{f_2}: p \rightarrow o_{i+2}$ be the appropriate reduced paths in \mathcal{T}_0 corresponding to the edges f_t in \mathcal{C}_0 and denote their labels by $w_0(\pi_{f_t}) =: w_t$ ($t = 1, 2$). Then $\pi_{f_1} l_1^{-1}: o_{i-1} \rightarrow o_i$, $l_1 l_2^{-1}: o_i \rightarrow o_{i+1}$ and $l_2 \pi_{f_2}: o_{i+1} \rightarrow o_{i+2}$ are the reduced paths in \mathcal{T} corresponding to the edges k_{i-1}, k_i, k_{i+1} in \mathcal{C} . The labels of these paths are $w_1 y_1^{-1}$, $y_1 y_2^{-1}$ and $y_2 w_2$, respectively. By the inductive assumption on \mathcal{C}_0 and \mathcal{T}_0 we have $[w_1]_H = [x_1]_H$, $[w_2]_H = [x_2^{-1}]_H$ and $\text{co}(\pi_1^Q(w_1)) \subseteq \text{co}(\pi_1^Q(x_1))$, $\text{co}(\pi_1^Q(w_2)) \subseteq \text{co}(\pi_1^Q(x_2^{-1}))$, which implies

$$[w_1 y_1^{-1}]_H = [x_1 y_1^{-1}]_H \text{ and } [y_2 w_2]_H = [y_2 x_2^{-1}]_H$$

as well as

$$\text{co}(\pi_1^Q(w_1 y_1^{-1})) \subseteq \text{co}(\pi_1^Q(x_1 y_1^{-1})) \text{ and } \text{co}(\pi_1^Q(y_2 w_2)) \subseteq \text{co}(\pi_1^Q(y_2 x_2^{-1})).$$

In combination with (4.9) and (4.11), (4.12), (4.13) we get for each one of the edges $k = k_{i-1}, k_i, k_{i+1}$ that

$$[w_k]_H = [v(k)]_H \text{ and } \text{co}(\pi_1^Q(w_k)) \subseteq \text{co}(\pi_1^Q(v(k)))$$

where π_k is the reduced path $\alpha k \rightarrow \omega k$ in \mathcal{T} and $w_k = w(\pi_k)$. For every edge k different from $k_{i-1}^{\pm 1}, k_i^{\pm 1}, k_{i+1}^{\pm 1}$ the claim follows from the inductive

hypothesis on \mathcal{C}_0 and \mathcal{T}_0 , since the corresponding path π_k runs entirely in \mathcal{T}_0 .

Figure 6 suggests that we have tacitly assumed that $\eta_i \neq \delta \neq \eta_{i+1}$. However, as for the case of cycles of length 3, the discussion above also covers the special cases $\eta_i = \delta$ or $\delta = \eta_{i+1}$: one of the paths Λ_i and the corresponding word y_i reduces to the trivial path and empty word, respectively. This is not in conflict with the construction since we allow the tree \mathcal{T} to have empty labels.

4.2.4. Arbitrary graphs. The general case will be handled by induction on the number l of (geometric) edges of the connected graph \mathcal{E} on n vertices. Base of the induction is $l = n - 1$, in which case \mathcal{E} is a tree. We may take $\mathcal{T} = \mathcal{E}$ and $w = v$. The idea is the same as for the proof of Lemma 3.7 and is borrowed from Section 7 of [6]. The proof is actually simpler: we fix the group H in advance, it works for all connected graphs on at most n vertices and need not be expanded in the inductive process when we add edges. We assume that \mathcal{E} has no loop edges.

So, let $l > n - 1$ and assume that the claim of Lemma 4.15 is true for all connected graphs \mathcal{D} (without loop edges) on n vertices with fewer than l geometric edges; let $\mathcal{E} = (V \cup K, \alpha, \omega, {}^{-1})$ be connected (without loop edges) with $|K^\bullet| = l$, $|V| = n$ and an H -commuting labelling $v: K \rightarrow \tilde{A}^*$ be given.

Recall that the E -generated group G involved in the definition of H is n -acyclic. Let $k^{\pm 1}$ be a geometric edge of \mathcal{E} that is contained in a cycle subgraph of \mathcal{E} and let $\mathcal{D} = \mathcal{E} \setminus \{k, k^{-1}\}$. The labelling $v \upharpoonright K(\mathcal{D})$ also commutes over H , and so the inductive assumption applies. So there exist a tree \mathcal{X} with $V(\mathcal{D}) \subseteq V(\mathcal{X})$ and a labelling $w_{\mathcal{X}}: K(\mathcal{X}) \rightarrow \tilde{A}^*$ such that:

- all vertices of \mathcal{X} not in \mathcal{D} have degree at least 3,
- if π_e is the reduced path $\alpha e \rightarrow \omega e$ in \mathcal{X} , then $[v(e)]_H = [w_{\mathcal{X}}(\pi_e)]_H$ and $\text{co}(\pi_1^Q(w_{\mathcal{X}}(\pi_e))) \subseteq \text{co}(\pi_1^Q(v(e)))$.

Note that the condition on the degree of the vertices guarantees that the diameter of \mathcal{X} is at most $n - 1$.

Choose edges k_1, \dots, k_t in \mathcal{D} such that $k_1 \cdots k_t: \omega k \rightarrow \alpha k$ is a simple path (that is, there are no repetitions of edges), then $kk_1 \cdots k_t$ is a simple closed path in \mathcal{E} . Since v is H -commuting, $[v(k)v(k_1 \cdots k_t)]_H = 1_H$. Denote by $\pi: \omega k \rightarrow \alpha k$ and $\pi_j: \alpha k_j \rightarrow \omega k_j$ ($j = 1, \dots, t$) the respective reduced paths in \mathcal{X} . Then π is the reduced path of $\pi_1 \cdots \pi_t$, whence $[w_{\mathcal{X}}(\pi)]_F = [w_{\mathcal{X}}(\pi_1 \cdots \pi_t)]_F$, and therefore $[w_{\mathcal{X}}(\pi)]_H = [w_{\mathcal{X}}(\pi_1 \cdots \pi_t)]_H$. Since $[w_{\mathcal{X}}(\pi_j)]_H = [v(k_j)]_H$, we get $[v(k)w_{\mathcal{X}}(\pi)]_H = [v(k)v(k_1 \cdots k_t)]_H = 1_H$.

We look more carefully at the path π . Assume that $\pi = f_1 \cdots f_s$, where the f_j are edges in \mathcal{X} :

$$\omega k = p_0 \xrightarrow{f_1} p_1 \xrightarrow{f_2} \cdots \xrightarrow{f_{s-1}} p_{s-1} \xrightarrow{f_s} p_s = \alpha k;$$

since the diameter of \mathcal{X} is at most $n - 1$, the length of π is also at most $n - 1$. We add to \mathcal{X} a new edge $f_0: \alpha k = \omega \pi \rightarrow \alpha \pi = \omega k$ and its inverse

f_0^{-1} and consider the closed path $f_0\pi = f_0f_1 \cdots f_s$ of length $s+1 \leq n$. Let \mathcal{C} be the cycle graph spanned by this closed path. We define a labelling $w_{\mathcal{C}}$ of \mathcal{K} by $w_{\mathcal{C}}(f_0) := v(k)$ and $w_{\mathcal{C}}(f_j) = w_{\mathcal{X}}(f_j)$ for $j \neq 0$ so that $w_{\mathcal{C}}$ provides a labelling of \mathcal{C} which commutes over H . According to Section 4.2.3, there exists a tree \mathcal{Y} on a vertex set $V(\mathcal{Y}) \supseteq V(\mathcal{C}) = \{p_0, \dots, p_s\}$ containing the vertices of \mathcal{C} , and a labelling $w_{\mathcal{Y}}$ of \mathcal{Y} such that, for each of the edges f_j , if $\rho_j: \alpha f_j \rightarrow \omega f_j$ denotes the reduced path in \mathcal{Y} , then

$$[w_{\mathcal{Y}}(\rho_j)]_H = [w_{\mathcal{C}}(f_j)]_H \text{ and } \text{co}(\pi_1^{\mathcal{Q}}(w_{\mathcal{Y}}(\rho_j))) \subseteq \text{co}(\pi_1^{\mathcal{Q}}(w_{\mathcal{C}}(f_j))). \quad (4.14)$$

We note that $w_{\mathcal{C}}(f_j) = w_{\mathcal{X}}(f_j)$ if $j \neq 0$ and $w_{\mathcal{C}}(f_0) = v(k)$. We may assume that \mathcal{Y} and \mathcal{X} have no edges in common and share no vertices other than those of \mathcal{C} , viz. p_0, \dots, p_s . Now define a tree \mathcal{Z} as

$$\mathcal{Z} = \mathcal{Y} \cup \mathcal{X} \setminus \{f_1^{\pm 1}, \dots, f_s^{\pm 1}\}$$

and a labelling $w_{\mathcal{Z}}$ of \mathcal{Z} by

$$w_{\mathcal{Z}}(f) = \begin{cases} w_{\mathcal{X}}(f) & \text{if } f \in \mathcal{X} \\ w_{\mathcal{Y}}(f) & \text{if } f \in \mathcal{Y}. \end{cases}$$

The situation is completely the same as in the proof of Lemma 3.7, cf. Figure 2. We get a picture of how the trees \mathcal{Y} and $\mathcal{X}_0, \dots, \mathcal{X}_s$ are merged to form \mathcal{Z} ; the subtrees $\mathcal{X}_0, \mathcal{X}_1, \dots, \mathcal{X}_s$ indicate the connected components of the forest $\mathcal{X} \setminus \{f_1^{\pm 1}, \dots, f_s^{\pm 1}\}$, \mathcal{X}_i being the component of p_i in this forest. We note that $V(\mathcal{E}) = V(\mathcal{D}) \subseteq V(\mathcal{X}) \subseteq V(\mathcal{Z})$. Let e be an edge of \mathcal{E} and let $\varsigma_e: \alpha e \rightarrow \omega e$ be the reduced path in \mathcal{Z} . If $e = k$ then $\varsigma_e = \varsigma_k$ runs completely in \mathcal{Y} and we have

$$[v(k)]_H = [w_{\mathcal{C}}(f_0)]_H = [w_{\mathcal{Y}}(\rho_0)]_H = [w_{\mathcal{Z}}(\varsigma_k)]_H$$

and

$$\text{co}(\pi_1^{\mathcal{Q}}(v(k))) = \text{co}(\pi_1^{\mathcal{Q}}(w_{\mathcal{C}}(f_0))) \supseteq \text{co}(\pi_1^{\mathcal{Q}}(w_{\mathcal{Y}}(\rho_0))) = \text{co}(\pi_1^{\mathcal{Q}}(w_{\mathcal{Z}}(\varsigma_k)))$$

by the construction of \mathcal{Y} from \mathcal{C} , see (4.14). Suppose that $e \neq k^{\pm 1}$, that is $e \in \mathcal{D}$. Then

$$[v(e)]_H = [w_{\mathcal{X}}(\pi_e)]_H \text{ and } \text{co}(\pi_1^{\mathcal{Q}}(v(e))) \supseteq \text{co}(\pi_1^{\mathcal{Q}}(w_{\mathcal{X}}(\pi_e))) \quad (4.15)$$

by the inductive hypothesis on \mathcal{X} , where π_e is the reduced path $\alpha e \rightarrow \omega e$ in \mathcal{X} . The path π_e may use some segment of the arc spanned by $f_1 \cdots f_s$, so

$$\pi_e = \pi_1 f_l \cdots f_m \pi_2 \text{ or } \pi_e = \pi_1 f_m^{-1} \cdots f_l^{-1} \pi_2 \quad (4.16)$$

for some reduced paths π_1, π_2 in $\mathcal{X} \setminus \{f_1^{\pm 1}, \dots, f_s^{\pm 1}\}$ and $l \leq m$. We handle the first case, the second one is completely analogous. Let $\rho: \alpha f_l \rightarrow \omega f_m$ be the reduced path in \mathcal{Y} . Then the reduced path $\varsigma_e: \alpha e \rightarrow \omega e$ in \mathcal{Z} is $\varsigma_e = \pi_1 \rho \pi_2$ and from the definition of $w_{\mathcal{Z}}$:

$$w_{\mathcal{Z}}(\varsigma_e) = w_{\mathcal{X}}(\pi_1) w_{\mathcal{Y}}(\rho) w_{\mathcal{X}}(\pi_2). \quad (4.17)$$

The path ρ is the reduced form of $\rho_l \cdots \rho_m$; on the one hand, this implies $[w_{\mathcal{Y}}(\rho)]_F = [w_{\mathcal{Y}}(\rho_l \cdots \rho_m)]_F$ so that $[w_{\mathcal{Y}}(\rho)]_H = [w_{\mathcal{Y}}(\rho_l \cdots \rho_m)]_H$. On the

other hand we get $\text{co}(\pi_1^{\mathcal{Q}}(w_{\mathcal{Y}}(\rho))) \subseteq \text{co}(\pi_1^{\mathcal{Q}}(w_{\mathcal{Y}}(\rho_l \cdots \rho_m)))$. From (4.14) we have

$$[w_{\mathcal{Y}}(\rho_j)]_H = [w_{\mathcal{X}}(f_j)]_H \text{ and } \text{co}(\pi_1^{\mathcal{Q}}(w_{\mathcal{Y}}(\rho_j))) \subseteq \text{co}(\pi_1^{\mathcal{Q}}(w_{\mathcal{X}}(f_j)))$$

for every $j \neq 0$. So from (4.17) and the structure of π_e in (4.16) we find $[w_{\mathcal{X}}(\pi_e)]_H = [w_{\mathcal{Z}}(\varsigma_e)]_H$ and $\text{co}(\pi_1^{\mathcal{Q}}(w_{\mathcal{X}}(\pi_e))) \supseteq \text{co}(\pi_1^{\mathcal{Q}}(w_{\mathcal{Z}}(\varsigma_e)))$. In combination with (4.15) we get for $e \neq k^{\pm 1}$ (what we had above already for $e = k^{\pm 1}$):

$$[v(e)]_H = [w_{\mathcal{Z}}(\varsigma_e)]_H \text{ and } \text{co}(\pi_1^{\mathcal{Q}}(v(e))) \supseteq \text{co}(\pi_1^{\mathcal{Q}}(w_{\mathcal{Z}}(\varsigma_e))). \quad (4.18)$$

We are *almost* done, but not yet completely: it may happen that the tree \mathcal{Z} contains vertices not in $V(\mathcal{E})$ of degree smaller than 3. This may occur as follows: some intermediary vertices in the path $f_1 \cdots f_s$ may be not in $V(\mathcal{E}) = V(\mathcal{D})$ and have degree 3 in \mathcal{X} ; these vertices may have degree 1 in \mathcal{Y} (in fact, they do have degree 1 as can be checked in the construction presented in Section 4.2.3). Every such vertex has degree 2 in \mathcal{Z} . To rectify this, consider all maximal arcs spanned by paths

$$q_0 \xrightarrow[e_1]{} q_1 \xrightarrow[e_2]{} q_2 \cdots q_{t-1} \xrightarrow[e_t]{} q_t \quad (4.19)$$

of consecutive edges e_1, \dots, e_t in \mathcal{Z} such that the vertices q_1, \dots, q_{t-1} have degree 2 and are not in $V(\mathcal{E})$ while each of q_0 and q_t is either in $V(\mathcal{E})$ or has degree at least 3. In order to get the required tree \mathcal{T} from \mathcal{Z} , we may proceed as we did earlier (proofs of Lemma 3.9 and Lemma 3.7): just contract the path (4.19) to a single edge, that is, replace every arc spanned by a path of the form (4.19) by the single edge $e: q_0 \longrightarrow q_t$ (and its inverse) and label e in \mathcal{T} by $w_{\mathcal{T}}(e) := w_{\mathcal{Z}}(e_1)w_{\mathcal{Z}}(e_2) \cdots w_{\mathcal{Z}}(e_t)$; edges e of \mathcal{Z} that are not on such a path remain unchanged in \mathcal{T} and we set $w_{\mathcal{T}}(e) := w_{\mathcal{Z}}(e)$. In the new tree \mathcal{T} just vertices from $V(\mathcal{E})$ may have degree less than 3. Every reduced path in \mathcal{Z} that traverses one of the edges e_i in (4.19) needs to traverse the entire arc (4.19). So, if $\varsigma_e: \alpha e \longrightarrow \omega e$ denotes the reduced path in \mathcal{Z} and $\tau_e: \alpha e \longrightarrow \omega e$ is the one in \mathcal{T} for an edge $e \in \mathcal{E}$, then $w_{\mathcal{Z}}(\varsigma_e) = w_{\mathcal{T}}(\tau_e)$ (i.e. the labellings of these paths coincide). Finally, from (4.18) we get

$$[v(e)]_H = [w_{\mathcal{T}}(\tau_e)]_H \text{ and } \text{co}(\pi_1^{\mathcal{Q}}(v(e))) \supseteq \text{co}(\pi_1^{\mathcal{Q}}(w_{\mathcal{T}}(\tau_e)))$$

for every edge $e \in \mathcal{E}$, where $\tau_e: \alpha e \longrightarrow \omega e$ denotes the reduced path in the new tree \mathcal{T} . \square

While no inverse monoid occurs explicitly in the preconditions and the formulation of Lemma 4.15 there is one “lurking from behind”. The precise meaning of this will become clear from Lemma 4.16 below, but intuitively it can be seen as follows. Let Q be an A -generated group (not necessarily finite) and consider the equivalence relation \sim_Q on the free involutory monoid \tilde{A}^* defined by

$$u \sim_Q v \iff [u]_Q = [v]_Q \text{ and } \text{co}(\pi_1^{\mathcal{Q}}(u)) = \text{co}(\pi_1^{\mathcal{Q}}(v)). \quad (4.20)$$

This means that two words u and v relate under \sim_Q if and only if

- (1) u and v evaluate equally in Q : $[u]_Q = [v]_Q =: g$ and

- (2) in the Cayley graph \mathcal{Q} of Q , the paths $1_Q \longrightarrow g$ labelled u and v , respectively, traverse the same geometric edges.

It turns out that \sim_Q is a congruence on the free involutory monoid \tilde{A}^* and the quotient $M(Q) := \tilde{A}^*/\sim_Q$ is an inverse monoid. The \sim_Q -class of a word $u \in \tilde{A}^*$ is determined by the pair $(\text{co}(\pi_1^Q(u)), [u]_Q) =: (R, g)$ where $R = \text{co}(\pi_1^Q(u)) \subseteq E = Q \times A$ and $g = [u]_Q \in Q$. Moreover, given two words, u and v , say, then, for $g = [u]_Q$,

$$\text{co}(\pi_1^Q(uv)) = \text{co}(\pi_1^Q(u)) \cup \text{co}(\pi_1^Q(v)) = \text{co}(\pi_1^Q(u)) \cup g \text{co}(\pi_1^Q(v)).$$

It follows that the pairs $(R, g) \in 2^E \times Q$ describing congruence classes modulo \sim_Q multiply by the rule $(R, g)(S, h) = (R \cup gS, gh)$. In other words, the inverse monoid $M(Q)$ can be realised as a submonoid of the semidirect product $2^E \rtimes Q$ where 2^E is the semilattice monoid (i.e. the commutative and idempotent monoid) of all subsets of E with \cup as binary operation and the action of Q on 2^E is the one induced by left multiplication of the elements of $E = Q \times A$ by elements of Q . One can check that in the inverse monoid $M(Q)$ the natural order is described by

$$(R, g) \leq (S, h) \text{ if and only if } g = h \text{ and } R \supseteq S.$$

Looking back at Lemma 4.15 we see that (in terms of the notation of that lemma) we *actually* have proved that the labelling w of the tree \mathcal{T} satisfies:

- (1) $[v(k)]_H = [w_k]_H$ and
- (2) $[v(k)]_{M(Q)} \leq [w_k]_{M(Q)}$.

Item (2) implies in particular that $[v(k)]_M \leq [w_k]_M$ for every A -generated inverse monoid M for which $M(Q) \twoheadrightarrow M$. The inverse monoid $M(Q)$ is called the *Margolis–Meakin-expansion* of Q [23]. It plays an important rôle in the theory of inverse monoids and will be of significant use in Section 5. Usually, the elements of $M(Q)$ are described by pairs (\mathcal{R}, g) where \mathcal{R} is a subgraph of \mathcal{Q} rather than pairs (R, g) with $R \in 2^E$. The connection between the pairs (\mathcal{R}, g) and (R, g) is clear, of course: instead of $R \subseteq E$ just take $\mathcal{R} := \langle R \rangle$, the subgraph of \mathcal{Q} spanned by R . This view allows one to characterise completely which pairs $(R, g) \in 2^E \times Q$ actually belong to $M(Q)$: (R, g) belongs to $M(Q)$ if and only if $\langle R \rangle$ is a finite connected subgraph of \mathcal{Q} containing 1_Q and g (with the natural convention that $\langle \emptyset \rangle = \{1_Q\}$). The collection of all A -generated inverse monoids M is partially ordered by $M \leq N$ if and only if $N \twoheadrightarrow M$. The special rôle of $M(Q)$ then lies in the fact that it is the maximum element of all A -generated inverse monoids M that are extended by Q . Essentially, this follows from the next lemma.

Let M be an A -generated inverse monoid and Q an A -generated group extending M ; recall the word content $\text{co}(w)$ of a word $w \in \tilde{E}^*$ according to Definition 4.2, for $E = Q \times A$.

Lemma 4.16. *For $u, v \in \tilde{A}^*$ such that $[u]_Q = [v]_Q$, the inclusion relationship $\langle \pi_1^Q(u) \rangle \subseteq \langle \pi_1^Q(v) \rangle$, or equivalently, $\text{co}(\pi_1^Q(u)) \subseteq \text{co}(\pi_1^Q(v))$, implies $[u]_M \geq [v]_M$.*

Proof. It is clear that $\langle \pi_1^Q(u) \rangle \subseteq \langle \pi_1^Q(v) \rangle$ if and only if $\text{co}(\pi_1^Q(u)) \subseteq \text{co}(\pi_1^Q(v))$. Let $u, v \in \tilde{A}^*$ be such that $[u]_Q = [v]_Q$ and $\langle \pi_1^Q(u) \rangle \subseteq \langle \pi_1^Q(v) \rangle$. We represent the A -generated inverse monoid M as a monoid of partial bijections on a finite set X ; we need to show that the partial bijection defined by v is a restriction of the one defined by u . Let \mathcal{P} be the A -graph on the vertex set X encoding the partial bijections $a \in \tilde{A}$. We consider the Q -compatible \tilde{A} -set $X_{\mathcal{P}}Q$ generated by X subject to \mathcal{P} defined in (2.2). Our assumption on Q implies that Q extends \mathcal{P} in the sense of Definition 2.3. Take any connected component \mathcal{P}_i of \mathcal{P} and let \mathcal{E}_i be the component of $X_{\mathcal{P}}Q$ containing \mathcal{P}_i as a subgraph. Then there exists, for every vertex y of \mathcal{E}_i , a canonical graph morphism $\varphi_y: \mathcal{Q} \rightarrow \mathcal{E}_i$ mapping 1_Q to y (\mathcal{Q} the Cayley graph of Q).⁷ Let $x \in X$ be a vertex of \mathcal{P}_i and consider the graph morphism φ_x . Since $\langle \pi_1^Q(v) \rangle \supseteq \langle \pi_1^Q(u) \rangle$ the same inclusion holds for the images of these graphs under φ_x . Hence, if the path in $\mathcal{E}_i \subseteq X_{\mathcal{P}}Q$ starting at x and being labelled v runs in \mathcal{P}_i , then so does the one labelled u . Consequently, if $x \cdot [v]_M$ is defined then so is $x \cdot [u]_M$. If both are defined then they are equal since

$$x \cdot [v]_M = x[v]_Q = x[u]_Q = x \cdot [u]_M.$$

It follows that $[v]_M \leq [u]_M$. \square

The preceding lemma is essentially folklore. Proofs of the lemma based on an analysis of the words u and v and not using the interpretation of M as a monoid of partial transformations can be found in [23, 37, 17, 13, 14], the present proof is sketched in [10] and perhaps elsewhere. From this lemma one can conclude that the canonical morphism $\tilde{A}^* \rightarrow M$ (for M extended by Q) factors through $M(Q)$, and so $M(Q)$ is indeed the maximum among all A -generated inverse monoids extended by Q . Therefore, Lemma 4.16 implies that if Q extends M then statement (2) $\text{co}(\pi_1^Q(v(k))) \supseteq \text{co}(\pi_1^Q(w_k))$ in Lemma 4.15 may be replaced by $[v(k)]_M \leq [w_k]_M$. We finally show that Lemma 4.15 implies Theorem 4.1.

Proof of Theorem 4.1 using Lemma 4.15 and Lemma 4.16. Let M and n be as in the statement of Theorem 4.1 and let \mathcal{E} be a connected graph on n vertices. Let \mathcal{D} be the graph obtained by removing all loop edges of \mathcal{E} . Let Q be a finite A -generated group extending M , H the A -generated group defined by (4.5) with G as in the statement of Lemma 4.15. By this lemma, there is a finite tree \mathcal{T} with $V(\mathcal{E}) = V(\mathcal{D}) \subseteq V(\mathcal{T})$ and a labelling $w: K(\mathcal{T}) \rightarrow \tilde{A}^*$ as in the statement of the lemma (for \mathcal{D} in the rôle of \mathcal{E}), but “ $\text{co}(\pi_1^Q(v(k))) \supseteq \text{co}(\pi_1^Q(w_k))$ ” replaced by “ $[v(k)]_M \leq [w_k]_M$ ”. Define the labelling $u: K(\mathcal{E}) \rightarrow$

⁷This follows from the fact that \mathcal{E}_i is isomorphic to the Schreier graph \mathcal{S}_H for some subgroup H of F containing the normal subgroup N of F for which $G = F/N$, and clearly $\mathcal{Q} \cong \mathcal{S}_N \rightarrow \mathcal{S}_H$.

\tilde{A}^* by $u(e) := v(e)v(e)^{-1}w_e$ for every edge e of \mathcal{E} where $w_e = w(\pi_e)$ and π_e is the reduced path $\alpha e \rightarrow \omega e$ in \mathcal{T} (which is empty if $\alpha e = \omega e$). For every loop edge e , π_e is an empty path and we have $[u(e)]_H = [v(e)]_H = 1_H$ since v commutes over H , and $[u(e)]_M = [v(e)]_M$ since $[v(e)]_Q = 1_Q$ implies that $[v(e)]_M$ is idempotent, hence $[v(e)]_M = [v(e)v(e)^{-1}]_M = [u(e)]_M$. For every non-loop edge the definition also implies $[u(e)]_H = [v(e)v(e)^{-1}w_e]_H = [v(e)]_H$ since $[v(e)]_H = [w_e]_H$ and $[u(e)]_M = [v(e)v(e)^{-1}w_e]_M = [v(e)]_M$ since $[v(e)]_M \leq [w_e]_M$. Moreover, u commutes over F : for every closed path $e_1 \cdots e_t$ in \mathcal{E} the reduced path of the corresponding product path $\pi_{e_1} \cdots \pi_{e_t}$ in \mathcal{T} is the empty path, hence the reduced form of the label of $e_1 \cdots e_t$ is the empty word. \square

4.3. Applications. We take Theorem 4.1 and see what we get if we apply it in the context of the different versions of the Ash/Herwig–Lascar/Ribes–Zalesskii Theorem.

4.3.1. Model-theoretic Herwig–Lascar. Looking at the proof in Section 3.3 we see that the Herwig–Lascar Theorem could also be formulated as follows.

Proposition 4.17. *Let σ be a finite relational signature, \mathfrak{X} a finite σ -structure on base set X endowed with a set \tilde{A} of partial automorphisms encoded in a finite A -graph \mathcal{P} on the vertex set X . Then for every finite set Σ of finite σ -structures there exists a finite A -generated group G extending \mathcal{P} which is such that, for every structure $\mathfrak{U} \in \Sigma$, for every homomorphism $\phi: \mathfrak{U} \rightarrow \mathfrak{X}_{\mathcal{P}}G$ there is a mapping $\psi: \phi(\mathfrak{U}) \rightarrow \mathfrak{X}_{\mathcal{P}}F$ such that the composite $\psi \circ \phi$ is a homomorphism from \mathfrak{U} to the free extension $\mathfrak{X}_{\mathcal{P}}F$ of \mathfrak{X} .*

So, in a sense, the proof in Section 3.3 of the Herwig–Lascar Theorem provides, for every finite σ -structure \mathfrak{X} with a set of partial automorphisms (encoded in an A -graph \mathcal{P}) and any finite collection Σ of finite σ -structures, a finite extension $\mathfrak{X}_{\mathcal{P}}G$, that — from point of view of Σ — *looks like* the free extension $\mathfrak{X}_{\mathcal{P}}F$. The group G , and hence the extension $\mathfrak{X}_{\mathcal{P}}G$, just depends on the maximum of the weights of the structures in Σ , not on their isomorphism types. A similar result in [29] uses criteria of hypergraph acyclicity in application to the hypergraph structure induced by the G -orbit of X in $\mathfrak{X}_{\mathcal{P}}G$. In the terminology of [29, Proposition 5.10] the hypergraph whose hyperedges are the images of the base set X under the automorphisms induced by G is *locally acyclic* in the sense that, up to a certain size bound, all its sub-hypergraphs are acyclic (i.e. chordal and conformal, or equivalently: tree-decomposable, see [15, 16] for the terminology).⁸ The following strengthening of the Herwig–Lascar Theorem lies in the fact that the resulting finite extension $\mathfrak{X}_{\mathcal{P}}G$ of the given relational structure \mathfrak{X} provides an even better approximation to the free extension $\mathfrak{X}_{\mathcal{P}}F$. In a sense, such extensions

⁸And Corollary 5.12 there uses this local similarity to the free extension $\mathfrak{X}_{\mathcal{P}}F$, which is globally tree-like in this sense, to rule out forbidden homomorphisms to establish EPPA for classes **Excl**(Σ) similar to a use of Theorem 4.18 below.

are *locally free*, see Figure 8 below. Other strengthenings of the Herwig–Lascar Theorem may be found in [20] and in some references cited in the latter. Recall the concepts used in Section 3.3: given a finite σ -structure $\mathfrak{X} = (X; (R^{\mathfrak{X}})_{R \in \sigma})$ together with a set \tilde{A} of partial automorphisms of \mathfrak{X} encoded in an A -graph \mathcal{P} on the vertex set X , $\mathfrak{X}_{\mathcal{P}}F$ is the free extension of \mathfrak{X} ; for any A -generated group G satisfying the conditions of Lemma 2.6 we have the extension $\mathfrak{X}_{\mathcal{P}}G$. Recall also the definition of the weight of a relational structure: $w(\mathfrak{U}) = |\text{rf}(\mathfrak{U})| + \sum_{R \in \sigma} |R^{\mathfrak{U}}|$, where $\text{rf}(\mathfrak{U})$ is the set of all elements of \mathfrak{U} not contained in any relational tuple of \mathfrak{U} . We denote the canonical morphism $F \twoheadrightarrow G$ by $f \mapsto \bar{f}$. Then the canonical mapping $\mathfrak{X}_{\mathcal{P}}F \twoheadrightarrow \mathfrak{X}_{\mathcal{P}}G$, $xf \mapsto x\bar{f}$ is a morphism of \tilde{A} -sets as well as of σ -structures.

Theorem 4.18. *Let $\mathfrak{X} = (X; (R^{\mathfrak{X}})_{R \in \sigma})$ be a finite σ -structure endowed with a set \tilde{A} of partial automorphisms encoded in an A -graph \mathcal{P} on the vertex set X . For every positive integer n there exists a finite A -generated group G for which the extension $\mathfrak{X}_{\mathcal{P}}G$ satisfies the following.*

For every weak substructure \mathfrak{U} of $\mathfrak{X}_{\mathcal{P}}G$ of weight $w(\mathfrak{U})$ at most n there is an isomorphism $\varphi: \mathfrak{U} \rightarrow \hat{\mathfrak{U}} \subseteq \mathfrak{X}_{\mathcal{P}}F$ onto a weak substructure $\hat{\mathfrak{U}}$ of $\mathfrak{X}_{\mathcal{P}}F$. In addition, the inverse isomorphism $\hat{\mathfrak{U}} \rightarrow \mathfrak{U}$ is induced by the canonical morphism $\mathfrak{X}_{\mathcal{P}}F \twoheadrightarrow \mathfrak{X}_{\mathcal{P}}G$.

Proof. We proceed as in the proof of Theorem 3.2 presented in Section 3.3. Let $M = \mathcal{T}(\mathcal{P})$ be the transition monoid of \mathcal{P} and G a finite A -generated group according to Theorem 4.1 for all connected graphs on at most n vertices. Let $\mathfrak{X}_{\mathcal{P}}G$ be the extension induced by G . There exists a subset $J \subseteq G$ of size at most n such that every relational tuple (t_1, \dots, t_r) of \mathfrak{U} and every element of \mathfrak{U} is contained in Xg for some $g \in J$. Define a graph structure \mathcal{J} on the set J exactly as in the proof of Theorem 3.2 given in Section 3.3, with the same labelling as in that proof. That is, the set of vertices $V(\mathcal{J})$ is J and two vertices g and h are connected by one or more edges if and only if $Xg \cap Xh \neq \emptyset$. More precisely, there will be $|Xg \cap Xh|$ many edges from g to h and also in the reverse direction. Suppose that $Xg \cap Xh = \{z_1, \dots, z_s\}$; then there are $x_1, \dots, x_s, y_1, \dots, y_s \in X$ such that $x_i g = z_i = y_i h$ for all i , hence $x_i g h^{-1} = y_i$ for all i . According to Lemma 2.2 there exist words $v_1, \dots, v_s \in \tilde{A}^*$ such that $x_i g h^{-1} = y_i = x_i \cdot [v_i]_M$ and $[v_i]_G = g h^{-1}$. In this situation we put edges $e_1, \dots, e_s: g \rightarrow h$ with labels v_1, \dots, v_s , respectively. The corresponding inverse edges $e_i^{-1}: h \rightarrow g$ then have labels v_i^{-1} and satisfy $y_i h g^{-1} = x_i = y_i \cdot [v_i^{-1}]_M$. We do this for all pairs $g, h \in J$ for which $Xg \cap Xh \neq \emptyset$. We have thus defined a word labelling of \mathcal{J} . This labelling is G -commuting since for every edge $e: g \rightarrow h$ the G -value of its label is $g h^{-1}$. Hence the label of any closed path has G -value 1_G . According to Theorem 4.1 there exists a relabelling which commutes over the free group F . Compared to the proof of Theorem 3.2 given in Section 3.3, the relabelling now satisfies an additional condition. Back there (that is, in the proof in Section 3.3) we already had $[v_i]_F = [w_i]_F = [w^{(g,h)}]_F$

and $[v_i]_M = [w_i]_M \leq [w^{(g,h)}]_M$ for the labels of $e_1, \dots, e_s: g \rightarrow h$ for specific $g, h \in G$ for which $Xg \cap Xh \neq \emptyset$. Now we additionally have that $[w^{(g,h)}]_G = [w_i]_G = [v_i]_G = gh^{-1}$. We proceed as in the earlier proof. Take a connected component of \mathcal{J} and choose a base vertex o . Choose any $f_o \in F$ such that $\overline{f_o} = o$; for every g in that component of \mathcal{J} let $f_g := [p]_F f_o$ where p is the label of some path $g \rightarrow o$ (the labelling is with respect to the latest relabelling which provided every edge $e: g \rightarrow h$ with the label $w^{(g,h)}$ mentioned above) — this does not depend on the choice of that path because of F -commutativity of the labelling. Taking into account that the G -value of the label of every path $g \rightarrow o$ is go^{-1} , it follows that $\overline{f_g} = [p]_F f_o = [p]_G \overline{f_o} = go^{-1}o = g$. Hence the choice of f_g satisfies $\overline{f_g} = g$. This holds for all $g \in J$ since this can be done for every connected component of \mathcal{J} . As in the proof of Theorem 3.2 in Section 3.3, the mapping $\mathfrak{U} \rightarrow \mathfrak{X}_{\mathcal{P}}F$ defined by

$$\begin{array}{ccccc} \mathfrak{U} \cap Xg & \longrightarrow & X & \longrightarrow & Xf_g \subseteq X_{\mathcal{P}}F \\ z & \longmapsto & zg^{-1} & \longmapsto & (zg^{-1})f_g \end{array} \quad (4.21)$$

is well-defined (independent of the choice of g) and a morphism of the σ -structure \mathfrak{U} . Since $\overline{f_g} = g$ for all g , it follows that the composition of this morphism with the canonical mapping $\mathfrak{X}_{\mathcal{P}}F \rightarrow \mathfrak{X}_{\mathcal{P}}G$ is the identity mapping on \mathfrak{U} . As a consequence, the mapping (4.21) $\mathfrak{U} \rightarrow \hat{\mathfrak{U}}$ is bijective and a homomorphism of σ -structures in both directions, hence an isomorphism. \square




FIGURE 8

We close with two remarks. (i) The mapping ψ in (4.21) is defined on the entire subset $Z := \bigcup_{g \in J} Xg$. Its composition with the canonical morphism

$\mathfrak{X}_{\mathcal{P}}F \rightarrow \mathfrak{X}_{\mathcal{P}}G$ is the identity mapping on Z . Hence ψ is a bijection $Z \rightarrow \psi(Z)$. While the inverse mapping ψ^{-1} is an injective homomorphism on the substructure of $\mathfrak{X}_{\mathcal{P}}F$ induced on $\psi(Z)$, the mapping ψ itself cannot be guaranteed to be a homomorphism defined on every weak substructure of $\mathfrak{X}_{\mathcal{P}}G$ on base set Z : it is just guaranteed to respect the relational tuples of Z that are completely contained in some of the translates Xg for some $g \in J$.

(ii) In the proof of Theorem 4.18 we had to choose, for the vertex o , an element f_o in F such that $\overline{f_o} = o$. There are many such choices and the set of all choices is given by $f_o N$ where $N \trianglelefteq F$ is the normal subgroup of F for which $F/N = G$. Two such choices f_o and f'_o are connected via $f'_o = f_o k$ for some $k \in N$. This leads to new choices for all elements $f'_g = [p]_F f'_o = [p]_F f_o k = f_g k$. So, for every $k \in N$, the substructure $\widehat{\mathfrak{U}}k$ of $\mathfrak{X}_{\mathcal{P}}F$ serves the same purpose as the original $\widehat{\mathfrak{U}}$. These substructures of $\mathfrak{X}_{\mathcal{P}}F$ are related by the automorphisms $y \mapsto yk$ for $k \in N$. Figure 8 gives an overview over the situation.

4.3.2. *Group-theoretic Herwig–Lascar.* Finally we formulate strengthened forms of the group-theoretic versions presented in Section 3.5. They are consequences of the following lemma.

Lemma 4.19. *Let $\mathcal{E} = (V \cup K; \alpha, \omega, {}^{-1})$ be a finite connected graph and $\ell: K \rightarrow 2^F$, $k \mapsto H_k g_k$ a coset labelling. Then there exists a finite A -generated group H such that, for every H -commuting word labelling $l: K \rightarrow F$ with $l(k) \in \ell(k)$ for all $k \in K$, there exists an F -commuting word labelling $l': K \rightarrow F$ with $l'(k) \in \ell(k)$ for all $k \in K$ such that, in addition, $[l(k)]_H = [l'(k)]_H$ for every $k \in K$.*

Remark 4.20. The statement without the last “in addition” condition follows from Ash’s Theorem. The strengthening consists entirely of the seemingly innocent additional condition saying that “[$l(k)$] $_H = [l'(k)]_H$ for every $k \in K$ ”, that is, that the two labellings are H -related.

Proof. For every geometric edge $e = k^\bullet \in K^\bullet$ let \mathcal{S}_e be the plain Stallings graph of the coset $H_k g_k$ and let $\mathcal{S} := \bigsqcup_{e \in K^\bullet} \mathcal{S}_e$ be the disjoint union of these graphs. Let H be a finite A -generated group as in Theorem 4.1 for $n = |V|$ and $M = \mathcal{T}(\mathcal{S})$. Let $l: K \rightarrow F$ be an H -commuting word labelling of \mathcal{E} with $l(k) \in \ell(k)$ for every k . Let $l': K \rightarrow F$ be an M - and H -related relabelling that commutes over F according to Theorem 4.1. We only need to argue that $l'(k) \in \ell(k)$ for every k . Now $l(k) \in \ell(k) = H_k g_k$ implies that $\iota \cdot [l(k)]_M = \tau$ holds in \mathcal{S}_{k^\bullet} where ι and τ are the initial and terminal vertices of the Stallings graph $\mathcal{S}_{H_k g_k} = (\iota, \mathcal{S}_{k^\bullet}, \tau)$. Since $[l(k)]_M = [l'(k)]_M$ it follows that $\iota \cdot [l'(k)]_M = \tau$ also holds in \mathcal{S}_{k^\bullet} which in turn implies that $l'(k) \in H_k g_k = \ell(k)$. \square

We are now ready to proceed as in Section 3.5. We replace right cosets with left cosets and get strengthened versions of Propositions 3.14, 3.18 and 3.20. We explicitly formulate the strengthened versions of the latter two.

Note that, in both cases, the strengthening entirely consists of the seemingly innocent additional condition that the solutions y_i belong to the cosets $x_i N$.

Proposition 4.21 (strengthened version of Proposition 3.18). *Let $\mathcal{E} = (V \cup K; \alpha, \omega, {}^{-1})$ be a finite graph and $\ell: V \cup K \rightarrow 2^F$, $p \mapsto g_p H_p$ an extended coset labelling and let $(X_p)_{p \in V \cup K}$ be variables. Then there exists a finite index normal subgroup N of F such that: if the system*

$$X_{\alpha k} X_k = X_{\omega k} \text{ for all } k \in K$$

has a solution $(x_p)_{p \in V \cup K}$ in $F^{V \cup K}$ subject to the constraints $X_p \in g_p H_p N$ for all $p \in V \cup K$, then it even has a solution $(y_p)_{p \in V \cup K}$ subject to the constraints $X_p \in g_p H_p$ (all $p \in V \cup K$), such that, in addition, $y_p \in x_p N$ for all $p \in V \cup K$. The normal subgroup N only depends on the size $|V|$ of V and the collection of cosets $\{g_p H_p : p \in V \cup K\}$.

Proposition 4.22 (strengthened version of Proposition 3.20). *Let I and L be finite index sets, $(X_i)_{i \in I}$ variables, $(g_l)_{l \in L}$ elements of F and $(H_l)_{l \in L}$ finitely generated subgroups of F ; let $S \subseteq I \times L$ and $T \subseteq I^2 \times L$. Then there exists a finite index normal subgroup N of F such that: if the system of constraints*

$$X_i \in g_l H_l N \text{ for all } (i, l) \in S \text{ and } X_j^{-1} X_i \in g_l H_l N \text{ for all } (j, i, l) \in T$$

has a solution $(x_i)_{i \in I} \in F^I$, then even the system

$$X_i \in g_l H_l \text{ for all } (i, l) \in S \text{ and } X_j^{-1} X_i \in g_l H_l \text{ for all } (j, i, l) \in T$$

has a solution $(y_i)_{i \in I} \in F^I$ such that, in addition, $y_i \in x_i N$ for every $i \in I$. The normal subgroup N only depends on the size $|I|$ of I and the collection of cosets $\{g_l H_l : l \in L\}$.

5. COMPARING THEOREM 3.3 WITH THEOREM 4.1

In view of Lemma 4.15 and the discussion after its proof, the essence of Theorem 4.1 may be summarised as in Proposition 5.1 below, seemingly in a purely group-theoretic manner and without the use of inverse monoids. Recall that for a finite A -generated group Q and a word x over \tilde{A} , $\pi_1^Q(x)$ denotes the path starting at 1_Q in the Cayley graph \mathcal{Q} of Q , considered as a word over $Q \times \tilde{A}$, and also recall the definition of the word content $\text{co}(_)$ of a word over $Q \times \tilde{A}$ according to Definition 4.2.

Proposition 5.1. *For every finite A -generated group Q and every positive integer n there exists a finite A -generated group H expanding Q such that the following holds. For every finite connected graph \mathcal{E} on at most n vertices and every H -commuting labelling $v: K(\mathcal{E}) \rightarrow \tilde{A}^*$ there exists a labelling $w: K(\mathcal{E}) \rightarrow \tilde{A}^*$ that commutes over F and such that for every edge k of \mathcal{E}*

- (1) $[v(k)]_H = [w(k)]_H$ and
- (2) $\text{co}(\pi_1^Q(v(k))) \supseteq \text{co}(\pi_1^Q(w(k)))$.

As we have seen, Theorem 4.1 is essentially the same as Proposition 5.1. The condition occurring in Lemma 4.15 concerning the absence of loop edges is of proof-technical nature and not relevant for the validity of the claim. It is interesting that the essence of Theorem 3.3 can be formulated in a similar inverse-monoid-free fashion. Here we refer to Definition 3.6, Lemma 3.7 and Proposition 3.22.

Proposition 5.2. *For every finite A -generated group Q and every positive integer n there exists a finite A -generated group H expanding Q such that the following holds. For every finite connected graph \mathcal{E} on at most n vertices and every H -commuting labelling $v: K(\mathcal{E}) \rightarrow \tilde{A}^*$ there exists a labelling $w: K(\mathcal{E}) \rightarrow \tilde{A}^*$ that commutes over F and such that for every edge k of \mathcal{E}*

- (1) $[v(k)]_Q = [w(k)]_Q$ and
- (2) $\text{co}(\pi_1^Q(v(k))) \supseteq \text{co}(\pi_1^Q(w(k)))$.

Again, the conditions on the absence of loop edges in \mathcal{E} in Definition 3.6 is of proof-technical nature and not relevant for the formulation of the essence. According to Ash's proof in [6] (cf. the discussion at the beginning of Section 3.6), the group H does not only depend on Q and the number of vertices of \mathcal{E} , but also on the cyclomatic number of \mathcal{E} . The proof in [7] of Proposition 3.22 by use of iterated expansions of the form $G \mapsto \mathbf{Ab}_m G$ also proceeds by induction on the maximal length of a cycle subgraph of \mathcal{E} and the cyclomatic number of \mathcal{E} . However, it also shows that *all* multiple edges can be handled by just one iteration. This implies that for any connected graph on n vertices at most $(n-1)\binom{n-1}{2} + 1$ iterations of expansions $G \mapsto \mathbf{Ab}_m G$ are needed to obtain the desired group H from Q (cf. the formulation of Proposition 3.22).

The essential (and only) difference between Proposition 5.1 and Proposition 5.2 obviously concerns item (1): while the group H of the stronger version guarantees that the old labelling v and the new labelling w are H -related, the weaker version only ensures that these labellings are Q -related. Actually, the series of expansions $Q \mapsto Q_1 \mapsto \cdots \mapsto Q_{N-1} \mapsto Q_N =: H$, where for every i , $Q_{i+1} = \mathbf{Ab}_{m_i} Q_i$ for some $m_i \geq 2$, leading to the expansion H of Proposition 5.2, even guarantees that the labellings v and w are Q_{N-1} -related. This is stronger than being Q -related, but still weaker than being Q_N -related, which would be required for a proof of Proposition 5.1 (a similar observation can be made for Ash's iterated expansions discussed in Section 3.6). In a sense, the series of expansions used to prove Proposition 5.2 always remains “one step away” from delivering a proof of Proposition 5.1. This behaviour has been observed by the second author in his doctoral dissertation [17]. Nevertheless, it could a priori still be possible that v and w are H -related for $H := Q_N$ for some sufficiently large N . The second author conjectured that this is not the case for any N . As a consequence of Theorem 5.4 below, this conjecture turns out to be true. This also provides some evidence that Theorem 4.1 is substantially stronger than Theorem 3.3.

We start with some definitions. For a finite A -generated group Q let $H(Q, n)$ be the A -generated expansion of Q defined in (4.5), with the intermediate $(Q \times A)$ -generated group G described in Section 4.1.2 being n -acyclic. Moreover recall, for an integer $m \geq 2$, the expansion $Q \mapsto \mathbf{Ab}_m Q$ introduced at the end of Section 3.6 in the context of Proposition 3.22. For integers $m_1, \dots, m_t \geq 2$ denote by $\mathbf{Ab}_{m_t} \cdots \mathbf{Ab}_{m_1} Q$ the expansion of Q obtained by iterated expansions $Q \mapsto \mathbf{Ab}_{m_1} Q \mapsto \mathbf{Ab}_{m_2}(\mathbf{Ab}_{m_1} Q) \mapsto \dots \mapsto \mathbf{Ab}_{m_t}(\dots(\mathbf{Ab}_{m_1} Q)\dots)$.

An *expansion series* **exps** is a mapping $Q \mapsto \mathbf{exps}(Q)$ that, for every A , assigns to every finite A -generated group Q a set $\mathbf{exps}(Q)$ of finite A -generated expansions of Q . Examples of expansion series that are significant in the present context are **solk** and **conc** where:

- i) **solk**(Q) = $\{\mathbf{Ab}_{m_t} \cdots \mathbf{Ab}_{m_1} Q : t \geq 1, m_1, \dots, m_t \geq 2\}$
- ii) **conc**(Q) = $\{H(Q, n) : n \geq 5\}$.

Here **solk** stands for *solvable kernel* while **conc** stands for *connected content*.

Definition 5.3. An expansion series **exps** is *sufficient* for Theorem 3.3 if, for every finite alphabet A , every A -generated group Q that extends an A -generated inverse monoid M admits, for every finite graph \mathcal{E} , an expansion $H \in \mathbf{exps}(Q)$ satisfying the requirements of Theorem 3.3. Sufficiency for Theorem 4.1 is analogously defined.

For example, the expansion series **conc** is sufficient for Theorem 4.1 (the proof of this claim is the content of Section 4.2) and hence is sufficient also for Theorem 3.3. Proposition 3.22 tells us that the expansion series **solk** is sufficient for Theorem 3.3. The main point of the present section is to show that **solk** is not sufficient for Theorem 4.1. This may be interpreted as saying that Theorem 4.1 is strictly stronger (or harder) than Theorem 3.3. Indeed, the expansion series **solk** provides a tool that enables one to prove Theorem 3.3 while it is not possible to prove Theorem 4.1 by use of this instrument.

Theorem 5.4. *The expansion series **solk** is sufficient for Theorem 3.3, but is not sufficient for Theorem 4.1.*

We start to recall some concepts and results from the theory of group varieties; a standard reference for this topic is the monograph [26]. For two varieties \mathbf{U} and \mathbf{V} , their product \mathbf{UV} consists of all groups G having a normal subgroup $N \in \mathbf{U}$ for which the quotient G/N belongs to \mathbf{V} . The class \mathbf{UV} is again a variety and for any three varieties $\mathbf{U}, \mathbf{V}, \mathbf{W}$ we have $(\mathbf{UV})\mathbf{W} = \mathbf{U}(\mathbf{VW})$. For a variety \mathbf{U} and a group G the \mathbf{U} -verbal subgroup of G denoted by $\mathbf{U}(G)$ is the (unique) smallest normal subgroup N of G for which G/N belongs to \mathbf{U} . The term “verbal” in this context stems from the fact that $\mathbf{U}(G)$ is comprised of all values of words w such that $w = 1$ is an identity of \mathbf{U} .⁹ For example, for the variety \mathbf{Ab} of all Abelian groups,

⁹Here “value of a word” has a slightly different meaning than in Section 2.

the \mathbf{Ab} -verbal subgroup of a group G is the subgroup of G generated by all *commutators*, that is, all elements of G of the form $ghg^{-1}h^{-1}$ for $g, h \in G$. In the context of formations of finite groups the term \mathbf{U} -*residual* is used instead of \mathbf{U} -verbal subgroup. For two varieties \mathbf{U} and \mathbf{V} and a group G we have that

$$\mathbf{U}(\mathbf{V}(G)) = (\mathbf{UV})(G). \quad (5.1)$$

Next, let $Q = F/N$ be a not necessarily finite A -generated group and \mathbf{U} be a variety of groups. The \mathbf{U} -*universal expansion* ${}^{\mathbf{U}}Q$ is defined to be

$${}^{\mathbf{U}}Q := F/\mathbf{U}(N). \quad (5.2)$$

The \mathbf{U} -universal expansion of Q is the largest A -generated expansion H of Q such that the kernel of the canonical morphism $H \twoheadrightarrow Q$ belongs to \mathbf{U} . That is, the canonical morphism ${}^{\mathbf{U}}Q \twoheadrightarrow Q$ factors through any such expansion H of Q with kernel in \mathbf{U} . For two varieties \mathbf{U} and \mathbf{V} and an A -generated group Q we have ${}^{\mathbf{U}}({}^{\mathbf{V}}Q) = {}^{\mathbf{UV}}Q$ — this follows immediately from (5.1). In particular, the notation of the iterated expansion $\mathbf{Ab}_{m_t} \cdots \mathbf{Ab}_{m_1} Q$ chosen at the beginning of the present section is consistent with this: $\mathbf{Ab}_{m_t} \cdots \mathbf{Ab}_{m_1} Q$ denotes the universal expansion of Q with respect to the product variety $\mathbf{Ab}_{m_t} \cdots \mathbf{Ab}_{m_1}$ where \mathbf{Ab}_{m_i} denotes the variety defined by the identities $xy = yx$ and $x^{m_i} = 1$, which consists of all Abelian groups of exponent dividing m_i . A variety \mathbf{U} is *locally finite* if every finitely generated member of \mathbf{U} is finite; the product \mathbf{UV} of two locally finite varieties \mathbf{U} and \mathbf{V} is again locally finite. In general, the \mathbf{U} -universal expansion ${}^{\mathbf{U}}Q$ need not be finite, even for finite Q . However, ${}^{\mathbf{U}}Q$ is finite for finite Q provided that \mathbf{U} is locally finite. We note that the varieties $\mathbf{Ab}_{m_t} \cdots \mathbf{Ab}_{m_1}$ we deal with are all locally finite.

In the following we give a description of groups of the form ${}^{\mathbf{U}}Q$, similar to the construction of the group H in (4.5). In this context we note that if Q happens to be the A -generated free object in the variety \mathbf{V} then ${}^{\mathbf{U}}Q$ is the A -generated free object in the variety \mathbf{UV} . This fact will be exploited later. So let Q be an A -generated group and $E = Q \times A$ the set of positive edges of the Cayley graph \mathcal{Q} of Q . Let G be the E -generated free object in \mathbf{U} . Now Q acts as a group of permutations on E via $(g, a) \mapsto {}^q(g, a) = (qg, a)$ for $q \in Q$ and $(g, a) \in E$. Since G is free in \mathbf{U} , this action extends to an action of Q on G by automorphisms on the left, denoted $\gamma \mapsto {}^q\gamma$ for $q \in Q$ and $\gamma \in G$, and we may form the semidirect product $G \rtimes Q$. Let H be the following A -generated subgroup of $G \rtimes Q$:

$$H := \langle [(1_Q, a)]_G, [a]_Q : a \in A \rangle \leq G \rtimes Q. \quad (5.3)$$

Just as in the context of (4.5), for a word x over \tilde{A} , we have $[x]_H = ([\pi_1^Q(x)]_G, [x]_Q)$ (see (2.4) in [14]). The following is folklore (see e.g. [11]).

Proposition 5.5. *The group H defined in (5.3) is a realisation of ${}^{\mathbf{U}}Q$. In particular, if Q is the A -generated free object in the variety \mathbf{V} , then H is a realisation of the A -generated free object in \mathbf{UV} .*

In Lemma 5.7 below we shall discuss two crucial identities, the first of which holds in the variety $\mathbf{Ab}_{m_t} \cdots \mathbf{Ab}_{m_1}$ while the second one fails. These identities will be the essential ingredient for proving the desired result. For elements w_1, \dots, w_n of an involutory monoid define $[w_1, \dots, w_n]$ inductively by $[w_1] := w_1$, $[w_1, w_2] := w_1 w_2 w_1^{-1} w_2^{-1}$ and $[w_1, \dots, w_n] := [[w_1, \dots, w_{n-1}], w_n]$. We first need the following lemma.

Lemma 5.6. *Let S, Q be groups, S abelian and such that Q acts on S by automorphisms on the left. Let $a = (h, g), b = (x, 1_Q) \in S \rtimes Q$. Then the following hold:*

- (1) $[a, b]$ is independent of h ,
- (2) $[a, b] \neq 1_{S \rtimes Q}$ if and only if ${}^g x \neq x$.

Proof. By direct calculation we get

$$\begin{aligned} aba^{-1}b^{-1} &= (h, g)(x, 1_Q)({}^{g^{-1}}h^{-1}, g^{-1})(x^{-1}, 1_Q) \\ &= (h {}^g x, g)({}^{g^{-1}}h^{-1} \cdot {}^{g^{-1}}x^{-1}, g^{-1}) \\ &= (h {}^g x h^{-1} x^{-1}, 1_Q) = ({}^g x x^{-1}, 1_Q) \end{aligned}$$

since S is abelian. We see that $({}^g x x^{-1}, 1_Q)$ is independent of h and coincides with the identity element $(1_S, 1_Q)$ if and only if ${}^g x = x$. \square

We come to the result about identities in $\mathbf{Ab}_{m_t} \cdots \mathbf{Ab}_{m_1}$ announced above. The first author is grateful to R. Bryant and A. Krasilnikov for having brought this crucial result to his attention back in the year 2002.

Lemma 5.7. *Let $t \geq 1$ and $m_1, \dots, m_t \geq 2$ and let H be the free object on three letters x, y, z in the variety $\mathbf{Ab}_{m_t} \cdots \mathbf{Ab}_{m_1}$. Then the following hold in H :*

- (1) $[xyx^{-1}, z^{m_1}, z^{m_1 m_2}, \dots, z^{m_1 \cdots m_{t-1}}] = [y, z^{m_1}, z^{m_1 m_2}, \dots, z^{m_1 \cdots m_{t-1}}]$
- (2) $[y, z^{m_1}, z^{m_1 m_2}, \dots, z^{m_1 \cdots m_{t-1}}] \neq 1_H$.

Proof. The proof is by induction on t . For $t = 1$ this reduces to $xyx^{-1} = y$ and $y \neq 1_H$, which is trivial since \mathbf{Ab}_{m_1} is Abelian and non-trivial (that is, $m_1 \geq 2$). So let $t \geq 2$ and suppose that the claim is true for $t - 1$. We realise H as in (5.3) for Q the $A = \{x, y, z\}$ -generated free object in $\mathbf{Ab}_{m_{t-1}} \cdots \mathbf{Ab}_{m_1}$ and G the $(Q \times A)$ -generated free object in \mathbf{Ab}_{m_t} . Let $w = [y, z^{m_1}, z^{m_1 m_2}, \dots, z^{m_1 \cdots m_{t-2}}]$, $w' = [xyx^{-1}, z^{m_1}, z^{m_1 m_2}, \dots, z^{m_1 \cdots m_{t-2}}]$ and $v = z^{m_1 \cdots m_{t-1}}$ considered as words over \tilde{A} . We evaluate these words in H . By the inductive assumption, w and w' evaluate identically in Q ; hence $a := [w]_H = (h, g)$ and $a' := [w']_H = (h', g)$ with $[w]_Q = g = [w']_Q$ for some $h, h' \in G$ and $g \in Q$. All members of the variety $\mathbf{Ab}_{m_{t-1}} \cdots \mathbf{Ab}_{m_1}$ have exponents dividing $m_1 \cdots m_{t-1}$, hence $[v]_Q = 1_Q$ so that $b := [v]_H = (k, 1_Q)$ for some $k \in G$. By Lemma 5.6 (1) we find that $[a, b] = [a', b]$ which proves assertion (1) of the Lemma.

In order to prove item (2) we need to show that ${}^g k \neq k$. For this we note that, by the inductive assumption, $[w]_Q = g \neq 1_Q$. It follows from the

structure of the word w that g is not in the subgroup of Q generated by z , or, in other words, the content of g is $\{y, z\}$. (Note that, since Q is free in $\mathbf{Ab}_{m_{t-1}} \cdots \mathbf{Ab}_{m_1}$, it is in particular retractable, and therefore has a content function.) On the other hand, k belongs to the subgroup of G generated by all elements of the form (q, z) with q in the subgroup of Q generated by z . More precisely,

$$k = [(1_Q, z)([z]_Q, z)([z^2]_Q, z) \cdots ([z^{m_1 \cdots m_{t-1}-1}]_Q, z)]_G.$$

Since the order of $[z]_Q$ is $m_1 \cdots m_{t-1}$, the elements $([z^i]_Q, z)$ are pairwise distinct, hence none of them cancel with each other, and in particular $k \neq 1_G$. Since g has content $\{y, z\}$, $[z^i]_Q \neq g[z^j]_Q$ for all i, j and the latter elements $g[z^j]_Q$ are also pairwise distinct. From this we get ${}^g k \neq k$. \square

In order to continue we recall the *Margolis–Meakin expansion* $M(Q)$ of an A -generated group Q introduced in Section 4.2, which may be described as follows. Recall that the group Q acts on its Cayley graph \mathcal{Q} by left multiplication. For a subgraph \mathcal{K} of \mathcal{Q} and $g \in Q$ we denote by ${}^g \mathcal{K}$ the image of \mathcal{K} under the action of g . For a given A -generated group Q , the Margolis–Meakin expansion $M(Q)$ consists of all pairs (\mathcal{K}, g) where $g \in Q$ and \mathcal{K} is a finite connected subgraph of the Cayley graph \mathcal{Q} of Q containing the vertices 1_Q and g . Endowed with the multiplication

$$(\mathcal{K}, g)(\mathcal{L}, h) = (\mathcal{K} \cup {}^g \mathcal{L}, gh)$$

and involution

$$(\mathcal{K}, g)^{-1} = ({}^{g^{-1}} \mathcal{K}, g^{-1})$$

the set $M(Q)$ becomes an A -generated inverse monoid with identity element $(\{1_Q\}, 1_Q)$, and with the assignment function

$$A \rightarrow M(Q), \quad a \mapsto (\langle (1_Q, a) \rangle, [a]_Q).$$

The value of the word $p \in \tilde{A}^*$ in $M(Q)$ is



$$[p]_{M(Q)} = (\langle \pi_1^{\mathcal{Q}}(p) \rangle, [p]_Q),$$

where $\pi_1^{\mathcal{Q}}(p)$ is the path in \mathcal{Q} starting at 1_Q and having label p ; the natural partial order on $M(Q)$ is given by

$$(\mathcal{K}, g) \leq (\mathcal{L}, h) \text{ if and only if } \mathcal{K} \supseteq \mathcal{L} \text{ and } g = h.$$

In the following we let $Q := \langle a, b : a^2 = b^2 = (ab)^2 = 1 \rangle$ be the Klein four-group and consider its Margolis–Meakin expansion $M(Q)$. Let \mathcal{E} be the graph consisting of two vertices $o \neq p$, two edges $e_1, e_2 : o \rightarrow p$ and their inverses. It is clear that Q extends $M(Q)$. We show that no expansion $H \in \mathbf{sol}\mathbf{k}(Q)$ satisfies the requirements of Theorem 4.1 for $M = M(Q)$ and \mathcal{E} .

Suppose to the contrary that for some $t \geq 1$, $m_1, \dots, m_t \geq 2$ the expansion $H := \mathbf{Ab}_{m_t} \cdots \mathbf{Ab}_{m_1} Q$ satisfies the requirements of Theorem 4.1. We denote the positive edges of the Cayley graph \mathcal{Q} by the letters in $B :=$

FIGURE 9. Cayley graph \mathcal{Q} FIGURE 10. The graphs \mathcal{A}_1 and \mathcal{A}_2

$\{c, d, e, f, g, h, i, j\}$ as indicated in Figure 9, on the right-hand side. For the letters in $C := \{x, y, z\}$ set

$$w(x, y, z) := [xyx^{-1}, z^{m_1}, z^{m_1 m_2}, \dots, z^{m_1 \dots m_{t-1}}],$$

considered as a word over \tilde{C} . Then the word $w(e, gh, dc)$ over \tilde{B} forms a path in \mathcal{Q} closed at 1_Q and so does the word $w(f^{-1}, gh, dc)$. Here $w(e, gh, dc)$ denotes the word obtained from $w(x, y, z)$ by substituting $x \mapsto e, y \mapsto gh, z \mapsto dc$ and analogously for $w(f^{-1}, gh, dc)$. Hence there exist words u_1, u_2 over \tilde{A} such that $w(e, gh, dc) = \pi_1^{\mathcal{Q}}(u_1)$ and $w(f^{-1}, gh, dc) = \pi_1^{\mathcal{Q}}(u_2)$. As a first consequence we have that $[u_1]_M = (\mathcal{A}_1, 1_Q)$ and $[u_2]_M = (\mathcal{A}_2, 1_Q)$ for \mathcal{A}_1 and \mathcal{A}_2 as indicated in Figure 10.

Recall that \mathcal{E} is the graph consisting of two vertices $o \neq p$, two edges $e_1, e_2: o \rightarrow p$ and their inverses and consider the word labelling of \mathcal{E} defined by $\ell(e_i) = u_i$. We represent H as in (5.3) with G the B -generated free object in $\mathbf{Ab}_{m_t} \cdots \mathbf{Ab}_{m_1}$. From Lemma 5.6 (1) it follows that in G we have

$$\begin{aligned} [\pi_1^{\mathcal{Q}}(u_1)]_G &= [w(e, gh, dc)]_G = [w(1, gh, dc)]_G \\ &= [w(f^{-1}, gh, dc)]_G = [\pi_1^{\mathcal{Q}}(u_2)]_G \end{aligned}$$

hence $[u_1]_H = [u_2]_H$ since clearly $[u_1]_Q = [u_2]_Q$. Consequently, ℓ commutes over H .

Suppose v_1, v_2 are words over \tilde{A} such that $\ell'(e_i) = v_i$ provides an F -commuting relabelling of ℓ that is M -related to ℓ . The latter means that $[u_i]_M = [v_i]_M$ and thus $[u_i]_M \leq [\text{red}(v_i)]_M$. Since the labelling ℓ' commutes over F we have $\text{red}(v_1) = \text{red}(v_2) =: v$ and therefore $[v]_M \geq [u_1]_M, [u_2]_M$. It follows that $\langle \pi_1^Q(v) \rangle \subseteq \mathcal{A}_1 \cap \mathcal{A}_2$. The latter inclusion is only possible if v is a power of b (positive or negative). If $[u_i]_H = [v_i]_H$ then also $[u_i]_H = [v]_H$ since v_1 and v_2 have the same reduced form. It follows from Lemma 5.6 (2) that $[w(1, gh, dc)]_G \neq 1_G$ and in particular

$$[w(1, gh, dc)]_G \neq 1_G = [w(1, 1, dc)]_G.$$

In other words, g and h belong to the content of $[w(1, gh, dc)]_G$, which is the first entry of $[u_1]_H$ as well as of $[u_2]_H$. On the other hand, if v is a power of b (positive or negative) then g and h cannot belong to the content of $[\pi_1^Q(v)]_G$, which is the first entry of $[v]_H$. Consequently, $[u_1]_H \neq [v]_H \neq [u_2]_H$, a contradiction. In fact, we have shown that for no expansion $H \in \text{sol}(Q)$ can the A -generated direct product $M(Q) \times H$ be an F -inverse cover of $M(Q)$, see [14, Section 2.7].

The above arguments apply to the Margolis–Meakin expansion $M(Q)$ of any finite A -generated group Q whose Cayley graph admits two disjoint cycle subgraphs having two arcs \mathcal{B}_1 and \mathcal{B}_2 connecting them such that the intersection $\mathcal{B}_1 \cap \mathcal{B}_2$ does not connect the two cycles.

REFERENCES

- [1] J. Almeida, Hyperdecidable pseudovarieties and the calculation of semidirect products, *Internat. J. Algebra Comput.* **9** (1999), no. 3–4, 241–261.
- [2] J. Almeida, Dynamics of implicit operations and tameness of pseudovarieties of groups, *Trans. Amer. Math. Soc.* **354** (2002), 387–411.
- [3] J. Almeida, J. C. Costa and M. Zeitoun, Complete reducibility of systems of equations with respect to \mathbf{R} , *Portugal. Math. (N.S.)* **64** (4) (2007), 445–508.
- [4] J. Almeida and M. Delgado, Sur certains systèmes d’équations avec contraintes dans un groupe libre, *Portugal. Math.* **56** (1999), no. 4, 409–417.
- [5] J. Almeida and M. Delgado, Sur certains systèmes d’équations avec contraintes dans un groupe libre — addenda, *Portugal. Math.* **58** (2001), no. 4, 379–387.
- [6] C. J. Ash, Inevitable graphs: a proof of the type II conjecture and some related decision procedures, *Internat. J. Algebra Comput.* **1** (1991), no. 1, 127–146.
- [7] K. Auinger, A new proof of the Rhodes type II conjecture, *Intern. J. Algebra Comput.* **14** (2004), no. 5–6, 551–568.
- [8] K. Auinger, The geometry of profinite graphs revisited, *Groups Geom. Dyn.* **11** (2017), no. 1, 139–164.
- [9] K. Auinger and B. Steinberg, The geometry of profinite graphs with applications to free groups and finite monoids, *Trans. Amer. Math. Soc.* **356** (2004), no. 2, 805–851.
- [10] K. Auinger and B. Steinberg, A constructive version of the Ribes–Zalesskii product theorem, *Math. Z.* **250** (2005), 287–297.
- [11] K. Auinger and M. B. Szendrei, On F -inverse covers of inverse monoids, *J. Pure Appl. Algebra* **204** (2006), 493–506.
- [12] K. Auinger and A. Bors, Formations of finite groups with the M. Hall property, *J. Comb. Algebra* **3** (2019), no. 2, 113–152.

- [13] K. Auinger, G. Kudryavtseva and M. B. Szendrei, F -inverse monoids as algebraic structures in enriched signature, *Indiana Univ. Math. J.* **70** (2021), no. 5, 2107–2131.
- [14] K. Auinger, J. Bitterlich and M. Otto, Finite approximation of free groups I: The F -inverse cover problem, (2023) arXiv:2208.03273.
- [15] C. Beeri, R. Fagin, D. Maier, and M. Yannakakis. On the desirability of acyclic database schemes. *Journal of the ACM*, 30, pp. 497–513, 1983.
- [16] C. Berge. *Graphs and Hypergraphs*. North-Holland, 1973.
- [17] J. Bitterlich, Investigations into the Universal Algebra of Hypergraph Coverings and Applications. Dissertation, Technische Universität Darmstadt, 2019.
- [18] M. Hall, Coset representations in free groups, *Trans. Amer. Math. Soc.* **67** (1949), 421–432.
- [19] B. Herwig and D. Lascar, Extending partial automorphisms and the profinite topology on free groups, *Trans. Amer. Math. Soc.* **352** (2000), no. 5, 1985–2021.
- [20] J. Hubička, M. Konečný and J. Nešetřil, All those EPPA classes (strengthening of the Herwig–Lascar Theorem), *Trans. Amer. Math. Soc.* **375** (2022), 7601–7667.
- [21] I. Kapovich and A. Myasnikov, Stallings foldings and subgroups of free groups, *J. Algebra* **248** (2002), no. 2, 608–668.
- [22] M. V. Lawson, *Inverse Semigroups. The Theory of Partial Symmetries*, World Scientific, River Edge, NJ, 1998.
- [23] 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.
- [24] S. W. Margolis and J.-É. Pin, Inverse semigroups and extensions of groups by semi-lattices, *J. Algebra* **110** (1987), 277–297.
- [25] M. Petrich, *Inverse Semigroups*, Wiley, New York, 1984.
- [26] H. Neumann, *Varieties of Groups*, Ergebnisse der Mathematik und ihrer Grenzgebiete 37, Springer Verlag, Heidelberg, 1967.
- [27] M. Otto, Highly acyclic groups, hypergraph covers, and the guarded fragment, *J. ACM* **59** (2012), no. 1, Art. 5, 40 pp.
- [28] M. Otto, Acyclicity in finite groups and groupoids, (2024), arXiv:1806.08664v6.
- [29] M. Otto, Amalgamation and symmetry: from local to global consistency in the finite, (2020), arXiv:1709.00031.
- [30] J.-É. Pin and C. Reutenauer, A conjecture on the Hall topology of a free group, *Bull. London Math. Soc.* **23** (1991), no. 4, 356–362.
- [31] J. Rhodes and B. Steinberg, *The q -theory of finite semigroups*, Springer Monogr. Math., Springer, New York, 2009.
- [32] L. Ribes and P. Zalesskii, On the profinite topology on a free group, *Bull. London Math. Soc.* **25** (1993), no. 1, 37–43.
- [33] L. Ribes and P. A. Zalesskii, The pro- p topology of a free group and algorithmic problems in semigroups *Internat. J. Algebra Comput.* **4** (1994), no. 3, 359–374.
- [34] J.-P. Serre, *Trees*. Translated from the French original by John Stillwell. Corrected 2nd printing of the 1980 English translation. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003.
- [35] J. R. Stallings, Topology of finite graphs, *Invent. Math.* **71** (1983), no. 3, 551–565.
- [36] J. B. Stephen, Presentation of inverse monoids, *J. Pure Appl. Algebra* **63** (1990), 81–112.
- [37] B. Steinberg, Inverse automata and profinite topologies on a free group, *J. Pure Appl. Algebra* **167**, (2002), 341–359.
- [38] S. You, The product separability of the generalized free product of cyclic groups, *J. London Math. Soc.* **56**, (1997), 91–103.

FAKULTÄT FÜR MATHEMATIK, UNIVERSITÄT WIEN, OSKAR-MORGENSTERN-PLATZ 1,
A-1090 WIEN, AUSTRIA

Email address: `karl.auinger@univie.ac.at`

Email address: `bitt.j@protonmail.com`

DEPARTMENT OF MATHEMATICS, TECHNISCHE UNIVERSITÄT DARMSTADT, SCHLOSS-
GARTENSTRASSE 7, D-64289 DARMSTADT, GERMANY

Email address: `otto@mathematik.tu-darmstadt.de`